DOCUMENT RESUME

ED 030 314 EM 007 245

By-Simmons, Robert S., Silberman, Harry F.

A Plan for Research Toward Computer-Aided Instruction With Natural English. Technical Memorandum.

Tamalpais Union High School District, Larkspur, Calif.

Report No-TM-3623

Pub Date 21 Aug 67

Note-79p.

EDRS Price MF-$0.50 HC-$4.05

Descriptors-Cognitive Ability, Cognitive Processes, Computational Linguistics, *Computer Assisted Instruction, Computer Programs, *English, Feedback, Instructional Materials, Instructional Technology, Language, Linguistics, Morphemes, Programed Instruction, Programing, *Semantics, Syntax

Identifiers-Protosynthex III

A research program was planned to develop a first, experimental computer-assisted instruction system that would permit interaction with students in a subset natural English. At the base of this system was a model of cognition that would represent the knowledge content of the material to be taught and the student's current knowledge of it. A comparison of the model of student current knowledge and the model of the material to be taught offered a basis for feeding back appropriate information to the student to move him toward the eventual training goal. The research was planned in two concurrent phases. The first developed language processing technology based on Photosynthex III. The second used tutorial studies to discover appropriate methods for training the students. The first appendix consists of material outlining the Cognitive Structure Model for Verbal Understanding, and the second is a Sample of the Minimal lesson Structure. (Author/GO)

ED0 30314

TM-3623

REPRINT

TECHNICAL MEMORANDUM

(TM Series)

A PLAN FOR RESEARCH TOWARD COMPUTER-AIDED INSTRUCTION WITH NATURAL ENGLISH

Robert S. Simmons
Harry F. Silberman
August 21, 1967

SYSTEM
DEVELOPMENT
CORPORATION
2500 COLORADO AVE.
SANTA MONICA
CALIFORNIA
90406

ED030314

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
OFFICE OF EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL OFFICE OF EDUCATION POSITION OR POLICY.

ABSTRACT

A research program is planned to develop a first experimental computer aided instruction system that will permit interaction with students in a subset of natural English. At the base of this system is a model of cognition that has a capability to represent the knowledge content of the material to be taught and of the student's current knowledge of it. A comparison of the model of student current knowledge and the model of the material to be taught offers a basis for feeding back appropriate information to the student to move him toward the eventual training goal. The research is planned in two concurrent phases. The first develops language processing technology based on Protosynthex III. The second uses tutorial studies to discover appropriate methods for training the students.

EM007245

CONTENTS

1. INTRODUCTION AND SUMMARY 5
    1.1 INTRODUCTION 5
    1.2 SUMMARY 6
2. DISCUSSION 7
3. OVERVIEW OF THE RESEARCH PLAN 10
4. PHASE I, DESIGN OF NATURAL LANGUAGE PROCESSOR FOR THE CAI SYSTEM 13
5. PHASE II, BEHAVIORAL STUDIES 27
6. RESEARCH PLAN 34
7. REFERENCES 37
APPENDIX I i
APPENDIX II ii

1. INTRODUCTION AND SUMMARY

1.1 INTRODUCTION

Despite important developments in the technology of computer-aided instruction (CAI) over the past several years, today's systems are still notably weak in (1) communicating with students and teachers in a language natural to humans, (2) diagnosing the details and causes of a student's shortcomings in performance, and (3) providing individualized remedial sequences appropriate to the student's needs. These weaknesses can eventually be remedied by research toward the development of machines that can "understand" the text of a subject material and the student's mastery of it and, as a result of this understanding, act like a tutor in detecting student shortcomings and providing responsive remedial material in natural language forms.

Such an automated tutor must be based on a cognitive model that can contain linguistic and semantic knowledge sufficient to decode and generate natural language strings. It must be rich in background knowledge of relations that hold among objects in the world in order that it may relate the facts, assertions, and relations of an instructional area to a wider range of knowledge and so "understand" the content area to be taught.

As a result of several years of research at SDC on natural language processing on the one hand and computer-aided instruction on the other, systems such as the Protosynthex I, II and III language processors and the PLANIT CAI system have been developed. These form a basis of research technology from which we propose to develop a first experimental CAI system that includes a cognitive model and a natural language capability to enable it to act more like a human tutor.

1.2 SUMMARY

Two concurrent and interacting phases of research are planned. The first is concerned with linguistic, semantic, and logical studies leading to the construction of the natural language system based on our Protosynthex III cognitive model. The second is behavioral research in the form of tutorial experiments that simulate proposed configurations of the complete CAI system to discover how best to accomplish instruction in a responsively interactive CAI environment.

Protosynthex III is rapidly nearing completion as a sophisticated natural language processor that can be said to understand a fair subset of English statement and question forms. It serves as a software basis for the first phase in which the CAI system is to be developed. Research in this first phase requires the development of improved question-answering programs for evaluating student performance, more sophisticated semantic analysis methods for understanding a wider variety of English forms, and a sentence- and question-generating capability to allow for communication with students. Steps in the Phase II research include lesson planning, test and quiz construction, tutorial studies, and finally a formal experiment to evaluate the effectiveness of the instructional approach that is developed.

2. DISCUSSION

Generally a computer-aided instruction system (CAI) is designed to present a student with a sequence of content materials that he is to learn. At periodic intervals the student is tested to discover how much of the material he understands. As a result of his performance on these tests he may be branched to remedial instruction material, continue the content sequence, skip portions of the sequence, or in the final analysis it may be decided that he has completed the content sequence that he was to be taught.

Present-day CAI systems appear to us as notably weak in (1) communicating with the student and teacher in a language natural to humans, (2) diagnosing the causes of the student's shortcomings in performance, and (3) providing individualized remedial sequences appropriate to the student's diagnosed needs. Each of these weaknesses represents a major CAI research area in its own right. We believe, however, that the kernel problem of all three is the need to develop working models of the processes humans use to understand natural languages.

Thus far, CAI systems are capable of manipulating language only as strings of characters without regard to any referential meaning that these strings of characters may have. Consequently, all steps in the learning sequence, all allowable responses of the student to questions, and all responses of the CAI system to the student's answers must be explicitly programmed in advance by the lesson designer. And the results are that (1) remedial sequences for anticipated errors are determined by the subjective judgment of the lesson designer rather than by an objective determination of the student's state of knowledge at the time, and (2) the system cannot handle unanticipated responses or queries by the student. On the other hand, a CAI system based on a model of language understanding could both determine its course of action dynamically according to the student's present state of knowledge as diagnosed objectively by the system itself, and respond appropriately to paraphrases, too-general or too-specific answers, requests for information, and other "unanticipated" student inputs. Such a CAI system would provide a student with the flexible, responsive teaching partner capable of teaching him the subject matter effectively and with thorough conceptual understanding.

To understand language a computer system must have an ability to represent events, and relations that hold among events, in the world as it is perceived by people. With this capability, a language processing model can be said to "know" or "understand" verbal meanings. Without it, the computer moves words as objects rather than as symbols.

The understanding capability required as a basis for an effective language processor could also compare a student's knowledge with that required by a training goal. With such understanding of wherein the student falls short, it may also be able to furnish the remedial material most appropriate to that student's shortcomings. At the very least such a language processor will enable the student and teacher to communicate naturally with and through the CAI system.

State of Research in Language Processing: Attempts to understand natural languages sufficiently well to enable the construction of language processor that can automatically translate, answer questions, write essays, etc., have seen frequent publication in the literature of the last decade. This work has been surveyed by Simmons [1965, 1966], Kuno [1966], and by Bobrow, Fraser & Quillian [1967]. These surveys agree in showing (1) that syntactic analysis by computer is reasonably well understood, though still inadequate and (2) that semantic analysis remains in its infancy as a formal discipline, although some programs manage to disentangle a limited set of semantic complexities in English statements. Approximately twenty more-or-less general-purpose language processors (mainly question-answering systems) have been programmed for computer operation. It is generally the case with these that their aspirations have been more grandly conceived than executed. Each of these systems has nevertheless been able to deal reasonably well with a small subset of natural English and to answer questions using fairly sophisticated logical calculi. The inescapable conclusion reflected in these surveys is that no adequate language processor exists today and that a great deal of research is still required before a system that can deal in a sophisticated way with a large subset of English can come into existence.

Several recent lines of research by Quillian [1966], Abelson [1965], and Simmons _et_ _al_. [1966, 1967] have introduced models of cognitive (knowledge) structure that may prove sufficient to model verbal understanding for important segments of natural language. Theoretical papers by Woods [1966], and Schwarcz [1967], and experimental work by Kellogg [1967a, 1967b] have tended to confirm the validity of the semantic and logical approaches taken by Quillian and Simmons. In each of these six approaches semantic and logical processing of language has been treated explicitly and each has showed a significant potential for answering questions phrased in nontrivial subsets of natural English. Our own work, described later in this proposal, promises during the present year to result in the first completely programmed language processing system that allows for communication and questioning in a significant subset of natural English and, furthermore, offers a sound basis for verbal understanding by computer via a cognitive model that explicates meanings of verbal events.

Such natural language processors as the above are still much too experimental for practical usefulness. For several years they will remain laboratory _curiosa_ demonstrating that language can be understood by computers although at great cost and with small efficiency. It is not too early, however, to juxtapose the line of natural language research with other advanced research areas such as CAI wherein eventual applications lie. It is our confident expectation that an experimental CAI system based on the concepts of verbal understanding that are used in natural language processors will provide an important enrichment of research ideas and developments in both fields.

The CAI system we propose to construct over the next two years is advanced in concept, giving a first indication of how both student and teacher can freely interact with a computer-aided instructional system using natural language;

however, in this period of time, it cannot become a finished product ready for actual applications in teaching situations. We propose, instead, a first experimental natural language CAI system that will be useful to establish and test principles of communication, diagnosis, and remedial response in a natural language environment.

3. OVERVIEW OF THE RESEARCH PLAN

The system we are proposing implies a radical departure from existing CAI systems. Its design grows directly from considerations of what is implied for a CAI if natural language is the primary means of communications among student, computer, and teacher. Two concurrent and interacting phases of research are planned; one concerns linguistic, semantic, and logical work required for developing a natural language system that can model the content area and the student's mastery of it and measure the differences between the two models. The other is behavioral research primarily in the form of tutorial experiments that simulate proposed configurations of the CAI system, to discover how best to present the material, and how to use differences between the two models to diagnose and remedy shortcomings in the student's knowledge and responsively shape his progress toward the training goal.

Requirements of the Natural Language Processor: If natural language is to be understood in any nontrival sense by a computer (i.e., if a computer is to accept English statements and questions, perform syntactic and semantic analyses, answer questions, paraphrase statements and/or generate statements or questions, all for a significant subset of English) there must exist some representation of knowledge of relations that generally hold among events in the world as perceived by humans. This representation may be conceived of as a cognitive model of some portion of the world. Among world events, there exist symbolic events such as words. The cognitive model, if it is to understand a natural language, must have the capability of representing these verbal objects, the syntactic relations that hold among them, and their mapping onto the cognitive events they stand for. This mapping from symbolic events of a language onto cognitive events is what is required of a semantic theory.

In addition to the first-level capability for transforming from a string of natural language symbols into the cognitive structure, a second-level capability for determining when one string is equivalent to or implied by another is required if the cognitive structure is to prove useful for answering questions or detecting meaning-preserving paraphrases. At this second level the system is required at minimum to have a capability for following deductive chains of reasoning. A model with both of these capabilities, developed by Simmons, Burger and Schwarcz, is described briefly in Section 2.3 and in more detail in Appendix I.

To the extent that such a system for understanding language can be used as a basis for an automated instructional system, it suggests a unique design based on its own capabilities for understanding the subject matter to be taught. The structure of information required by the model to understand and use language also has the capability to represent an understanding of the content area to be taught. A similar cognitive structure based on student responses to quizzes can represent the student's knowledge of the subject area at any given instant. A comparison of the two models at any stage in the instructional process should show what the student has achieved or what he lacks, and it may also imply the sequence in which information is to be presented and misinformation corrected to further his progress toward the goal.

The CAI approach using natural language can be seen in the following proposed training sequence.

1. Use the language processing system with human help to produce the initial model (C1) of content to be taught.

2. Pretest the student's knowledge of the subject area with a series of questions whose answers form the basis for the first model (S1) of the student's achievement.

3. Compare model S1 with C1 and choose short or long course depending on the size of the discrepancy.

4. Assign a unit of textual material (in either case) and follow this by a short quiz to test the student's mastery of the unit.
5. Augment the student model S1 with the content structure of the essay material he generates in 4.
6. Compare S1 with the portion of C1 relevant to the text unit of 4 to discover gaps, wrong connections, and irrelevancies. Generate questions or statements as remedial material to correct the student's knowledge as represented by C1.
   - 6a. Generate a set of questions testing only the remedial material and repreat steps 5 and 6 until discrepancies between S1 and C1 decrease to an acceptable level.
7. Iterate steps 4, 5 and 6 until student has met the criterion for mastery of the entire content area.

The natural language system at the base of the proposed CAI system is described in Section III of the proposal. This system will be completed, enlarged, and modified to fit the CAI system requirements. A special modification of the question-answering system will be developed to compare the student model as though it were a question to the content model, and to return gaps, errors and irrelevancies. A system will be designed and programmed to generate meaningful English statements and questions based on earlier experimentation by Simmons & Londe [1965] and Klein & Simmons [1964]. A control system embodying the algorithms for selecting and presenting remedial material will be programmed following findings from behavioral studies which are described as a second phase of the project. Repeated experimental runs with the resulting complete CAI language processor will be performed to accumulate sufficient bodies of linguistic, semantic and logical rules to enable it to understand and respond to the student's actual training sequences. All of the proposed developments of the language processor will be oriented to the subset of English chosen for an area in physiology that will be selected as experimental instructional material.

Phase II Behavioral Studies: The result of the comparison of the student model to the content model in various iterations will show discrepancies in the form of knowledge gaps, errors of fact, and irrelevancies. Just how this information should best be used to control the system's generation of remedial material can best be discovered initially by behavioral experimentation in tutorial studies. A tutorial study is one in which the experimentor simulates the CAI system, leads the student through the same training sequence as the system would, but uses his own understanding as an educator or tutor to discover where the student is having difficulty and how best to correct it.

In the present case the tutorial experiment will be designed to discover how the measured discrepancies between student and content models can be used best to generate remedial material to correct the discrepancies.

The outcome of this line of experimentation will be the design basis of algorithms that control the CAI system's generation of remedial material for individuals on the basis of their performance. Further trial and revision cycles will be conducted on-line with the prototype CAI system. A formal experiment will then be conducted to assess the effect of the machine's natural language capability on student learning.

4. PHASE I, DESIGN OF NATURAL LANGUAGE PROCESSOR FOR THE CAI SYSTEM

At the base of the proposed CAI system, a sophisticated natural language processor is required. The language processor must accept and model an understanding of text, student questions and answers, and generate questions and statements in response to student communications. It must also be able to model the student's knowledge of the content area and compare this model to its content model. In addition to the content of the instructional material, the language processing system requires additional information in the form of general facts, inference rules, semantic meaning postulates, etc., in order to deal with it and the student in an understanding and responsive fashion.

In order to accomplish these requirements a model of cognitive structure and a complicated chain of programs including semantic and syntactic analysis and logical inference capabilities make up the language processing system. The bulk of these programs have been developed over the past several years in the System Development Corporation's Synthex project and can be modified as necessary for the CAI system. These programs and the required modifications are described briefly below and in more detail in Appendix I.

The Cognitive Model: The essential requirement of a language processing system is that it be able to represent the meaning of words, sentences, and larger units of text. In our view meaning is attached to a sensation or a symbol by embedding it in a network of relations among other perceptions. The cognitive model in order to represent meanings must be able to model relations that hold among objects in the world as perceived by humans. The representation of meanings and knowledge is what we mean by the term, cognitive structure.

Our model of cognitive structure derives from a theory of structure proposed by Allport [1955] in the psychological context of theories of perception. The primitive elements of our model are events and relations. An event is defined recursively either as an object or as an event-relation-event (E-R-E) triple. A relation is defined in extension as the set of pairs of events that it connects; intensionally a relation can be defined by having a set of properties such as transivity, reflexivity, etc., each having associated rules of reference.

Any perception, fact or happening, no matter how complex, can be conceived as a single event that can be expanded into a nested structure of (E-R-E) triples.* The entire structure of a person's knowledge at the cognitive or conceptual level can thus be expressed as a single event; or at the base of the nervous system, the excitation of two connected neurons is also an event that may be described at a deeper level as molecular events in relation to other molecular events.

---

* From a logician's point of view the **E-R-E** structure can be seen as a structure of binary relations of the form R(**E,E**) and this statement is equivalent to the logician's assertion that any event can be described in a formal language.

Our interest is not in describing neural, molecular or atomic events; instead we wish to be able to model the objects and relations of textual materials. Not surprisingly the structure of English can be resolved to this same E-R-E format. A language string such as:

The condor is a large vulture,

can be syntactically analyzed

((condor art the) is ((vulture adj large) art a)).

Elementary linguistic events in this structure are exemplified by "condor," "the," "vulture," "large," and "a." Terms such as "art(icle)" "is," and "adj(ective)" are taken as linguistic relation terms. Complex events include the entire sentence and the triples (condor art the), (vulture adj large), etc.

Events at the linguistic level, however, are ambiguous with respect to possible meaning. For example, "condor" might refer to a bird, an airplane or perhaps a game played by children. "Is" may denote the relation of equality, of superset or of attribution--among others. In the cognitive model an event should have an unambiguous representation. It is the task of semantic analysis to map the English words of a linguistic structure into an unambiguous set of cognitive events. For the example sentence, the cognitive or formal representation is:

((condor$_1$ Q generic) SUP (vulture$_1$ SIZE large$_1$) Q specific))[*]

---

[*] Equivalent to the formalization in functional calculus:

$$\forall\ \text{condor}_1 \Big[[\text{condor}_1 \subset \text{condor}] \supset \exists\ \text{vulture}_1 \Big[[\text{vulture}_1 \subset \text{vulture}] \wedge \text{SIZE}\ (\text{vulture}_1, \text{large})\Big] \wedge [\text{condor}_1 \subset \text{vulture}_1]\Big]$$

Q stands for quantifier, SUP and SIZE for superset and size relations respectively. The subscripts signify particular tokens of condor, vulture and large.

This formal structure of the cognitive model can be expressed as a directed graph with labeled nodes as follows:

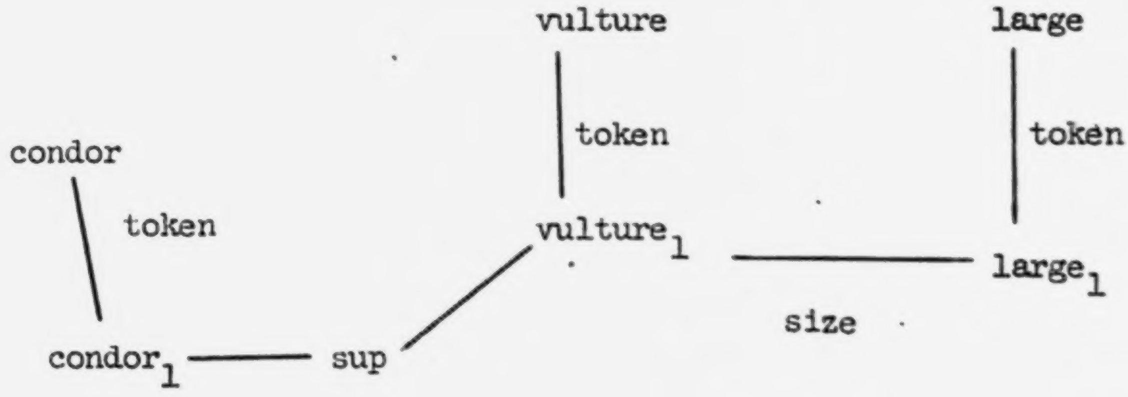

Meaning in this structure is derived from the interrelations of events and by the properties* attached to the relations that connect events. For example when we add the notion that vulture is a bird, the structure is expanded by the addition of (vulture SUP bird) and meaning is thereby added to the structure. When we know that the relation SUP is characterized by the properties transitivity, reflexivity, and antisymmetry, the meaningful inferences that condor is a bird, condor is a condor, vulture is a vulture, vulture is not a condor, etc., are implied by deductive inference rules associated with these properties.

This brief description of the cognitive structure model leaves many questions unanswered. Most significant among these include the procedure for assigning relational terms such as SUP, SUB, EQUIV, SIZE, LOC, etc., and the use of quantifiers as signified by "each," "all," "every," "some," "the," "a(n)," etc., in English. Both of these are very difficult problems that are currently receiving attention from linguists and logicians both on our project and elsewhere. Some further discussion of the model and of these and other problems associated with it is included in Appendix I. We propose to direct research toward formalizing the vocabulary of relational terms and better reflecting the quantificational structure of natural language.

---

* For our system, a property is essentially a rule of inference (see p. 19 for examples).

The use of the model in answering questions and drawing inferences will be described following brief statements of the syntactic and semantic analysis programs that are used to convert from English sentences or questions into the structure of the model.

Syntactic Analyzer: The object of syntactic analysis in the language processor is to transform a complex sentence such as the following:

The condor of North America, called the California Condor, is the largest land bird on the continent,

into a set of simple nested triplets such as those below:

((((condor art the) of (America N* North)) called (((condor N California) art the) is ((((bird N land) adj largest) art the) on (continent art the)))) .

These nested triples can be arranged in the tree structure on page 14 and labeled to show their correspondence with the usual phrase structure analysis presented by linguists. The parentheses simply preserve the tree structure.

The syntactic analysis or parsing is obtained from an SDC-developed parsing system called PLP-II, which accumulates its grammar as a result of being told word-class and dependency information about each of the sentences it experiences. The system is described in more detail in Appendix I and in Burger, Long and Simmons [1966]. It is a well-developed system, programmed in LISP, and it has the capability to deal with a considerable range of complex English sentences. Recently, a limited capability for finding the antecedents of pronouns and pronominal adjectives has been added to the system. This latter capability is based on work currently in progress by Olney and Londe in their anaphoric and discourse analysis research [Olney 1965].

---

* N for noun; modification by a noun is a linguistic relation.

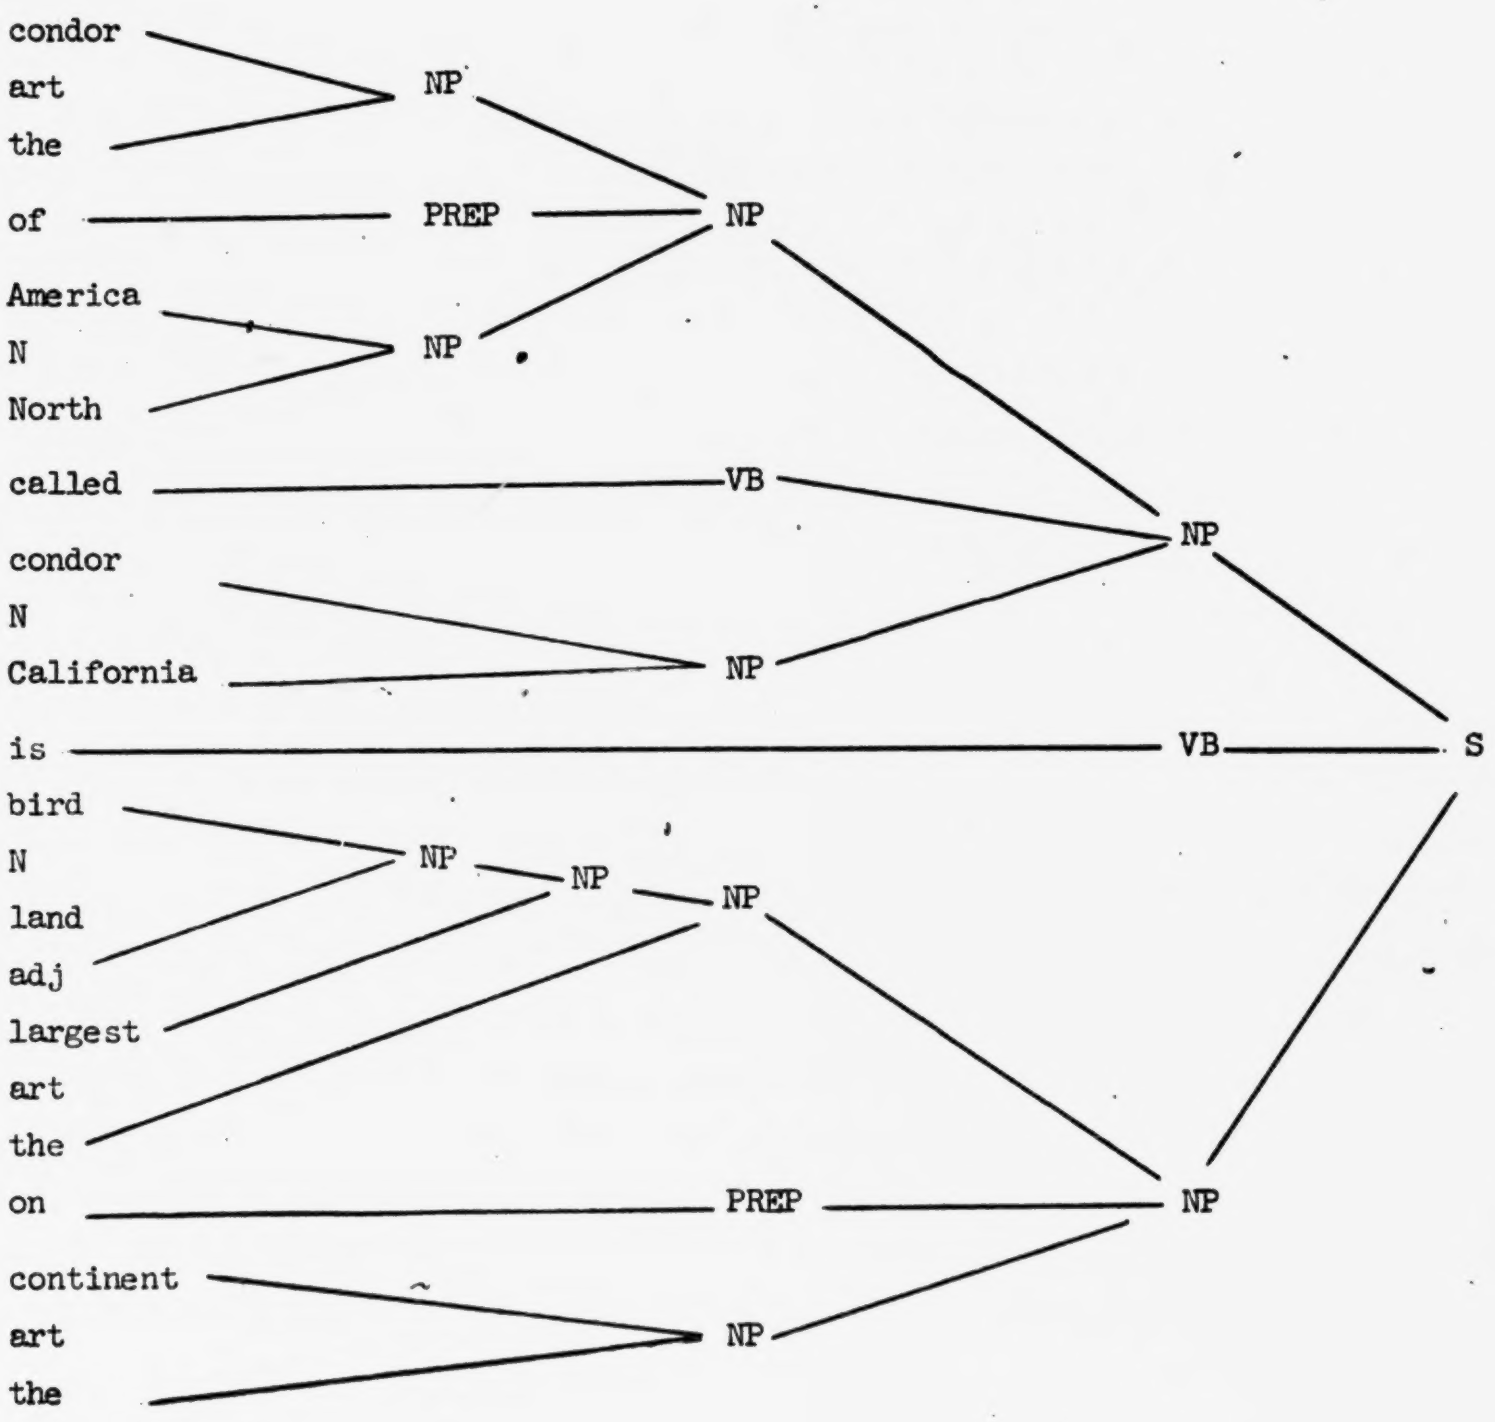
condor
art
the
NP
of
PREP
NP
America
N
North
NP
called
VB
NP
condor
N
California
NP
is
VB
S
bird
N
NP
land
adj
NP
largest
NP
art
the
on
PREP
NP
continent
art
the
NP

A Syntactic Structure of Nested Triples

Semantic Analysis: Our first version of a semantic analyzer is only now being programmed. Probably before June 1967, several LISP versions of this system will have been built and experimented with since the programs are not intrinsically complicated.

Semantic analysis for our language processor is defined as mapping from the ambiguous English words of a syntactic kernel (of the type obtained from PLP-II) into unambiguous objects of the cognitive model. For example, the syntactic kernel (pitcher hit batter) is several ways ambiguous. It might mean that a baseball player hit a baseball batter, that a glass pitcher dipped into a liquid batter, that the glass pitcher hit a man or even conceivably that a man hit the liquid batter. In the cognitive model however, each of these interpretations, if valid, must be represented as an unequivocal relation between two unambiguous cognitive objects. The cognitive model provides unique objects to represent the pitcher that is a container, the batter that is a liquid, the man-batter and the man-pitcher. The task of the semantic analysis is to select appropriate objects onto which it can map the meaning of the words of the English kernel.

The example (pitcher hit batter) is truly ambiguous since no further context is offered. If we deal with the more complete context ((pitcher adj angry) struck (batter adj careless)) we, as persons, recognize that the ambiguities have been eliminated. The semantic system must provide this capability by recognizing that "angry" cannot ordinarily modify a container-type pitcher nor "careless" a liquid.

This task of mapping from ambiguous verbal symbols to unambiguous cognitive objects is accomplished with a simple highly interactive LISP program that uses meaning postulates* and subset-superset relational chains in the following way. Assuming that the program has not yet accumulated any relational structure and has no meaning postulates, the program takes its input as follows:

---

* A meaning postulate is defined by Carnap [1956] as a rule that states as much about the meaning of a term as is required for analyticity in the framework of a semantic system. In the present usage, a meaning postulate explicates elements of semantic meaning implied by English words.

((pitcher adj angry) struck (batter adj careless))
((man emot emot) hit (man attitude attitude))*

The program then builds the following structure of subset (SUB) and superset (SUP) relations:

pitcher	SUP	man
adj	SUB	emotion
angry	SUP	emotion
struck	SUB	hit
batter	SUP	man
adj	SUB	attitude
careless	SUP	attitude

In addition it constructs the following meaning postulates:

man	emot	emot
man	attitude	attitude
man	hit	man .

This is a sufficient structure to disambiguate the sentence in the following manner. Assume now that the system has these and other data and that, in the form of its nested syntactic triples, the sentence was given to the semantic analyzer as follows:

((pitcher adj angry) struck (batter adj careless))

The analyzer looks up the supersets for pitcher and it might discover pitcher SUP man and pitcher SUP container. It looks up the subset of adjective and discovers "adj SUB emotion" and "adj SUB attitude." For angry it might find only "angry SUP emotion." It then attempts to translate the words of the kernel to superset and subset terms as follows:

---

* The operator chose these particular terms by knowing the nature of the cognitive model. Essentially for each word in the syntactic kernels he asked either what is the superset or subset term in the system that encompasses the sense in which this word is being used.

(pitcher adj angry) = (man emotion emotion)
(man attitude emotion)
(container attitude emotion)
(container emotion emotion) .

At this point it tests each of the possible translations against its list of meaning postulates where it might discover the following:

(man emotion emotion)
(man attitude attitude) .

Intersection of the two lists shows that only the interpretation (man emotion emotion) corresponds to a meaning postulate. Consequently an unambiguous cognitive object--that token of pitcher that has the superset man--can be selected. In a similar manner the ambiguities of batter and of struck are eliminated, and finally the sentence is rendered unambiguous in interpreting that the pitcher-man hit the batter-man.

It can be inferred from this brief description that a meaning postulate is essentially a rule of inductive inference, the complete set of which outlines the system's knowledge of possible relations among objects. In attempting semantic analysis, contexts are translated to meaning postulate form, triple by triple. In nested triples such as ((pitcher emot angry) struck (batter attitude careless)) the heads of the more deeply nested triples are used for the translation; thus (pitcher struck batter) is the topmost triple of the example sentence, and in this case the terms pitcher and batter have already been resolved to unambiguous nodes in the cognitive model. With the aid of an additional meaning postulate (man hit man) and the relation (struck SUB hit), the ambiguities of this sentence are resolved. It may easily happen that two or more interpretations are legitimate at lower levels of nesting and these may (or may not) be resolved at the highest level of the sentence.

The proposed method has been tested extensively on small samples of text and it is expected that the programs will require little effort before they are suitable for inclusion in the language processor. Two extensions of our semantic research are contemplated; first, we expect to extend our bracketing to include sentence and paragraph units and our meaning postulates to include reasonable sentence sequences. The hoped-for effect will be to add a further degree of disambiguation by using larger contexts than single sentences offer. The second extension of the semantic research offers, instead of the two-stage syntactic-semantic process, the exciting prospect of transforming directly from strings of natural language into the bracketed unambiguous form of the cognitive model. Preliminary studies have suggested that this is a feasible approach using semantic classes instead of syntactic ones, and meaning postulates in place of phrase structure rules.

Questions that have not been adequately answered in this semantic approach include requirements concerning the size of data structure, the number of meaning postulates, and the choice of level at which the meaning postulates are phrased (i.e., for (aardvark has tongue) we might state meaning postulates as (mammal hasprt appendage) or (animal hasprt part) etc.). We have also given some consideration to the underlying theory of semantic structures that guides our work and have seen a continuity that relates it to the Katz-Fodor theory of semantic markers on the one hand and to Sparck Jones' essentially statistical theory of semantic classification on the other. This continuity of theory must be developed further.

Answering English Questions: After having transformed English text and questions into the formal structure of the cognitive model, question-answering resolves either to a simple matching procedure or (more often) to a process of using inference rules to transform equivalent data structures into the form of the question. The first case is illustrated below:

	Question:	(dormouse size little)
(a)	Data:	(dormouse size little)
	Answer:	yes

An example of the more complicated case follows:

Question: (dormouse LOC Europe)

(b) Data: ((dormouse size little) SUP (native LOC Europe))

Inference rule for complex product ((SUP C/P LOC) implies LOC)

Answer: yes

The most general form of inference rules used in the system is a nested triple with variables (X1, X2, etc., R1, R2, etc.) for which elements of the question or the data structure may be substituted. For example the rule for the right-collapsible property can be expressed as follows:

(c) $(((X_1\ R_1\ X_2) \wedge (X_2\ R_2\ X_3))\ \text{IMPLIES}\ (X1\ R_2\ X3))$

Since the relation SUP is right-collapsible in example (b) we could substitute data for variables in (c) as follows:

(((dormouse SUP native) ∧ (native LOC Europe)) implies (dormouse LOC Europe)) and find the implied statement (dormouse LOC Europe) corresponds to the form of the question. All inference rules used in the system could be expressed in such a form, but for those that have been found to be of frequent utility the more efficient procedure of programming them as primitive system operations has been used to produce substantial savings in operating time. The present inference maker uses the following set of properties and functions as system operations:

1. Symmetry (SYMM): a relation with this property has itself as its inverse, i.e., XRY → YRX.
2. Intersection (S*AND): (R1 S*AND R2) holds between X and Y if and only if both R1 and R2 hold between X and Y i.e., (XR1Y) and (XR2Y).
3. Complex Product (C/P): (R1 C/P R2) holds between X and Y if and only if there is some Z such that both (X R1 Z) and (Z R2 Y) hold.
4. Transitivity (TRANS): If a relation R has this property, then (R C/P R) implies R (for any X and Y).

5. Left Collapsibility (LCOL): If a relation R2 has this property, then R1 C/P R2) implies R1 for any relation R1.
6. Right Collapsibility (RCOL): If a relation R1 has this property, then (R2 C/P R1) implies R1 for any relation R2.

The most interesting finding that we have so far achieved from the work on answering questions is that, after the question and the text have been put into the formal cognitive structure, the problem of question-answering is analogous to one of theorem proving or general problem solving as studied by Newell & Simon [1963], Wang [1960], and others. The question is equivalent to a theorem, and data in the structure are analogous to axioms, general theorems, and other special theorems that have already been shown to be valid. The operation of question-answering is one of applying various legitimate transformations in the form of inference rules to the true data theorems to determine if some combination of them is equivalent to the question theorem.

This finding, while encouraging in that it places question-answering into a well structured field of logical inquiry, is disquieting in that it leaves no room for doubting that in large data structures, difficult questions may take considerable lengths of time for answering. The problem with a large data base is similar to the chess problem; although algorithms for finding a best move (answer) may exist, the possibility space of the chessboard (large data base) is so great that the time requirement may approach the indefinitely large.

Our most significant problem in question-answering is one of arranging the application of inference rules to minimize the possibility space to be searched for answers. This area must continue to receive a large share of our research effort. The use of inference rules in our question-answering system is discussed in more detail in Appendix I and in Simmons, Burger & Long [1966].

The CAI System: From the outline just presented it can be seen that the natural language processing component for the CAI system is, in its first experimental form, fairly complete. How it can be used to model the content of a training area and the knowledge of a student and how it can generate and present remedial materials remain to be discussed.

The proposed training sequence presented earlier (on pp. 6, 7) requires that, as a first step, a model be prepared for the content area of the text. To accomplish this an experimenter will type in the text, sentence by sentence, on a live, time-shared-computer teletype, furnishing syntactic and semantic information as required by the program system. This procedure will allow the language processor to accumulate its store of relevant linguistic and semantic information. The successful accomplishment of this stage for the training text automatically results in a cognitive model representing its content.

The experimenters will also prepare a set of questions that completely encompass the text and the points to be taught. A final examination will be prepared in alternate forms, one to be administered at the start of the sequence, the other at its terminus. Quizzes for each lesson segment of the text will also be prepared by the experimenters. All of these questions will be administered as inputs to the question-answering portion of the language processor and the experimenters will augment the content model with appropriate background information, linguistic and semantic information, and rules of inference until the system can answer all the questions.

The capability for modeling a student's knowledge of the content area is developed in a similar fashion based on his short essay answers to the questions. The procedure is planned as follows. A first experimental group of students will be instructed on the limitations of vocabulary and sentence structure that the language processor imposes. They will then be given the teaching sequence of examination-text-quiz, etc., and the terminal examination. Their answers to the examination questions will serve primarily as a basis for the experimenters to further augment the system's linguistic background and inference material so that it can better understand "unexpected" responses by students. The student's answers to the examination questions are modeled by the system in the same fashion that it modeled the text content. After an appropriate number of such shakedown experiments, the system should have accumulated enough knowledge of the types of responses to be expected from students to be ready for an experimental use as a training device. (At this point a major risk is that the accumulation of data may be so great as to preclude efficient use of storage devices.)

Comparing student response models to the content model can be seen as a form of question-answering in which the student model is treated as a question. A program will be prepared that will compare the two models, using a limited range of inference rules, and present discrepancies as a list of gaps (omitted information), incorrect facts, and irrelevant portions of the student's model.

Based on initial outcomes of Phase II experimentation, a program will be prepared to present this information back to the student. No mention has so far been made in this proposal of methods for generating English statements or questions. In fact, we have not yet attempted to do so from the present model. However in previous experimentation of this type by Klein & Simmons [1963] and Simmons and Londe [1964], and related work by Weizenbaum [1966] and Colby [1967], we have discovered that generating uncomplicated sentences is a comparatively straightforward process. We believe that a generation program for the present language processor can be developed in a short period of time, providing we restrict its capabilities to a very simple syntax and perhaps tolerate some stereotypy in its sentence patterns.

A critical point in the proposed line of research is the choice of methods to be used by the CAI system to respond to student model discrepancies with the generation of appropriate remedial material. The Phase II line of research, concurrent with Phase I, is described in the following section as an approach toward discovering optimal feedback and training strategies that will provide algorithms to control the generation and presentation of remedial material.

5. PHASE II, BEHAVIORAL STUDIES

Most existing CAI systems contain limited answer-processing routines that correct spelling errors in the student's responses, provide a keyword matching feature, and consider response latency in evaluating the student's answers. They typically discriminate only two aspects of the student's behavior--whether or not the student has made a response, and whether the response was correct or incorrect. A successful tutor, on the other hand, is capable of much finer discriminations. Consider two incorrect answers to the same question: one answer may indicate a complete lack of understanding by the student, while the other shows that the student understands all but one or two minor elements of the problem. The CAI system should react differently to these two incorrect answers. An answer revealing great lack of understanding might cause the machine to repeat a large part of its instruction through a lengthy remedial dialogue with the student, while an answer indicating almost complete comprehension might cause the machine to provide two or three appropriate hints sufficient to fill the gap in the student's knowledge.

An interactive CAI system would thus afford the student much more initiative in guiding the instruction, by shifting the emphasis away from the *tabula rasa* concept whereby the preplanned lesson is written into the student, and moving toward a natural language conversation with the student.

The objectives of the natural language CAI system will not be concerned with the mere learning of rote facts. In the proposed project, a much more highly complex skill will be established involving chains of verbal discourse leading to the solution of a problem whose answers are not available from a simple inspection of the textual material. At first, such interchange is overtly mediated by the natural language processing capability of the computer. As this process of verbal discourse becomes internalized with extended use of such instruction it is anticipated that the generalized problem-solving skills of the student will be improved.

Tutorial Studies: Tutorial studies simulating proposed configurations of the CAI system will be designed to discover how the measured discrepancies between student and content models can best be used to generate remedial material to correct the discrepancies.*

Within the tutorial sessions, the experimenters will try various tutoring strategies and note those that appear to be successful and could be implemented on a computer. Such strategies or algorithms will be included in the program system. Once the CAI system has been programmed, the lesson can be tried out with students, evaluated and revised. In evaluating and revising the CAI lesson it may be necessary to change the strategies used on various frames. This implies that the strategies or algorithms must be programmed in such a way that they may be applied to or denied to any frame.

The organization of the minimal lesson structure will be in frames, based on an analysis of the subject matter. A frame consists of a block of text followed by a question or problem (see Appendix II for examples). The lesson designer will specify the content of each frame. The machine (tutor) will analyze student answers (which may include questions as well as statements) and will generate a statement and/or question in reply to the student's response. The student will respond in turn to these machine-generated messages and the cycle will continue until a predetermined criterion of understanding is met, signaling the machine to move on to another frame.

The instructional logic to be used by the machine in generating the remedial dialogue must be determined empirically by tutorial sessions with individual students. Those tutoring strategies that result in effective learning will be abstracted and used in the design of the prototype machine. Some tutoring strategies that might be effective are:

1. Selective Reflection

   The student's response (if not completely acceptable) is reproduced for the student with erroneous or missing parts indicated by blanks.

---

*For examples of tutorial studies see Silberman and Coulson [1964], Coulson [1964], Newmark [1964], Silberman [1964] and Melaragno [1964].

Example:

TUTOR: Name the two types of retinal receptor cells.
STUDENT: Rods and fovea.
TUTOR: Rods and ________.
STUDENT: Rods and retina.
TUTOR: Rods and ________.
STUDENT: Rods and cones.

A more natural interactive variant of the same strategy might be:

TUTOR: Name the two types of retinal receptor cells.
STUDENT: Rods and fovea.
TUTOR: You are partially right. Rods are receptor cells, but fovea are not. What is another retinal receptor cell?
STUDENT: Retina.
TUTOR: No, not retina. What is another type of retinal receptor cell?
STUDENT: Cones.
...

2. Selective Repetition

Pertinent parts of the instructional material are repeated to the student if his answer is incomplete or wrong.

Example:

TUTOR: Name two types of retinal receptor cells.
STUDENT: Rods and fovea.
TUTOR: Well--remember we said: Two types of retinal receptor cells are rods and cones. Now, name two types of retinal receptor cells.
STUDENT: Rods and cones.
...

3. Sequencing of Algorithms

If the student's response involves more than one error or more than one kind of error, which error should be dealt with first?

Example:

TUTOR: How do the two types of end organs differ in innervation?

STUDENT: Rods have individual innervation.

TUTOR: But what about the other types of end organ?

STUDENT: Oh, cones don't have individual innervation.

TUTOR: You have it backwards. You have reversed the characteristics of rods and cones. Let's try again. How do the two types of end organs differ in innervation?

STUDENT: Rods do not have individual innervation; cones do.

...

Here the gap was attended to before the wrong response was corrected, but the wrong response could have been corrected first:

TUTOR: How do the two types of end organs differ in innervation?

STUDENT: Rods have individual innervation.

TUTOR: Not so. Rods do not have individual innervation, for several of these receptors are found to connect with multiple dendrites of a common bipolar cell. Now, how do the two types of end organs differ in innervation?

STUDENT: Rods do not have individual innervation.

TUTOR: Right. But what about cones?

STUDENT: Cones do have individual innervation.

4. Frame Termination Criteria

The good tutor knows when to drop one subject and pick up another. Either he detects the student's understanding of the point or he realizes that a new approach is called for. Examples of the former are the ends of each example above (the student has exhibited the desired behavior for the frame); examples of the latter are difficult to construct and such criteria must be discovered empirically.

There are many other questions whose answers may be discovered through tutorial study. Some are: How long should the messages be? Should the messages be posed in question or statement form, or both? If a student's answer contained both gaps and irrelevancies, which should be corrected first? How should synonyms be employed? To what extent is it necessary to employ meaningful words that appear with high frequency in the student's speaking vocabulary? How much redundancy should be provided in the information given to the student? How much use should be made of previews and summaries? Is it better to use forward or backward chaining procedures to teach the content most efficiently? What kind of reinforcers should be used? Presumably, the detailed analysis of the student's constructed responses by this system will in itself serve as a powerful reinforcer. How can the material be made more interesting? Should the instructional sequence be adapted to the student's own self evaluation? Should the sequence be responsive to how long it takes the student to respond to questions as well as to his pattern of errors?

The sample of questions listed above represents but a small subset of the population of questions that need to be answered in order to build an efficient instructional logic. Although it will not be possible to obtain answers to all these questions within the scope of this project, it will be possible to discover at least some effective instructional strategies although many strategies may exist. This will be accomplished by the tutorial studies using a succession of evaluation revision cycles on individual students. This work will tentatively use a three-hour sequence on physiology as subject matter content. Once an effective instructional sequence has been developed, a formal experiment will be conducted to determine the extent to which the effectiveness of this sequence is attributable to the capability of the natural language CAI system to analyze the unique responses of each learner. It is predicted that the performance of students receiving an instructional sequence that is uniquely tailored to their particular learning needs, as assessed by an analysis of their verbal behavior, will be superior with respect to scores on a criterion posttest to that of students receiving a sequence that is not so tailored.

Formal Experiment: The purpose of this experiment is to determine the extent to which the instructional effectiveness of the system that results from the tutorial studies can be attributed to the capability of the system to analyze student's questions and statements. A question-and-answering CAI program would allow the student to engage in a natural discourse with the computer, similar to a dialogue between student and tutor. The program would read English questions and text that are composed by the student. It would perform a grammatical and semantic analysis of the student's response, make appropriate inferences about the student's understanding of the concept to be taught, and generate questions and statements that are designed to enhance the student's understanding. It is proposed that a study be conducted to determine whether the question-answering program increases the effectiveness of computer-based instruction. It is hypothesized that the performance of students receiving an instructional sequence contingent on a detailed analysis of their responses will be superior with respect to scores on a criterion posttest to that of students receiving an instructional sequence that has not been tailored to their response.

Approximately 30 students will be selected from local colleges and universities in the Los Angeles area. One treatment group will be designated the responsive group and the other, the nonresponsive group. Members of the responsive group will receive sequences of information and questions determined by a detailed analysis of the responses they make during the teaching session. Throughout the instructional session, the machine will select an appropriate sequence of instructional messages for each student based upon his particular errors, e.g., incorrect information, conceptual gaps, and irrelevant information. In this way, the student is given only that instruction which he needs. Each member of the nonresponsive group will be paired at random with one member of the responsive group. The unique sequence of frames that each student in the responsive group generates will be presented to his mate in the nonresponsive group. Thus, pairs of students in the two groups will receive identical instructional sequences, but the machine will be responsive to the particular kind of errors made by students in the responsive group, and not necessarily responsive to errors made by students in the other group. Knowledge of results will consist of a simple statement of the correct answer to each question and will be the same for both groups. Every student will be told the following: "You will receive a sequence of instruction. The instruction will consist of a series of messages or frames. Each frame will consist of some information followed by one or more questions. Sometimes you will only receive questions with no accompanying information. You are to give any answer that you think appropriate. You are also free to ask any questions. Sometimes I will not be able to answer your questions and will tell you so. After the instructional session, you will receive a test."

No time restrictions will be placed on students during either the training period or the test period. Analysis of variance techniques will be used to analyze the test data.

6. RESEARCH PLAN

The research program is divided into two concurrent and interacting phases with subphases as outlined below and illustrated in Figure 1:

Phase I, Development of CAI Software

1. Assembly of first complete language processor based on present development of Protosynthex III.
2. Design and development of advanced semantic analyzer that transforms directly from strings of language into the formal cognitive structure.
3. Design and programming of a sentence generator to produce meaningful English statements and questions from the cognitive model.
4. Assembly of programs for initial version of the CAI system. Modifications to this system continue throughout the two-year period.
5. Modification of the question-answering logic to allow it to compare model of student knowledge with model of content.
6. Programming of algorithm to control sentence and question generation for remedial feedback to the student following findings from Phase II, subphase 5.
7. Development of content model on CAI system including the amassing of semantic and background information for understanding text and student language.
8. Shakedown trials with the CAI system to further its ability to deal with student language, questions, etc.
9. Instructional trials devoted to training students in the three-hour content sequence.

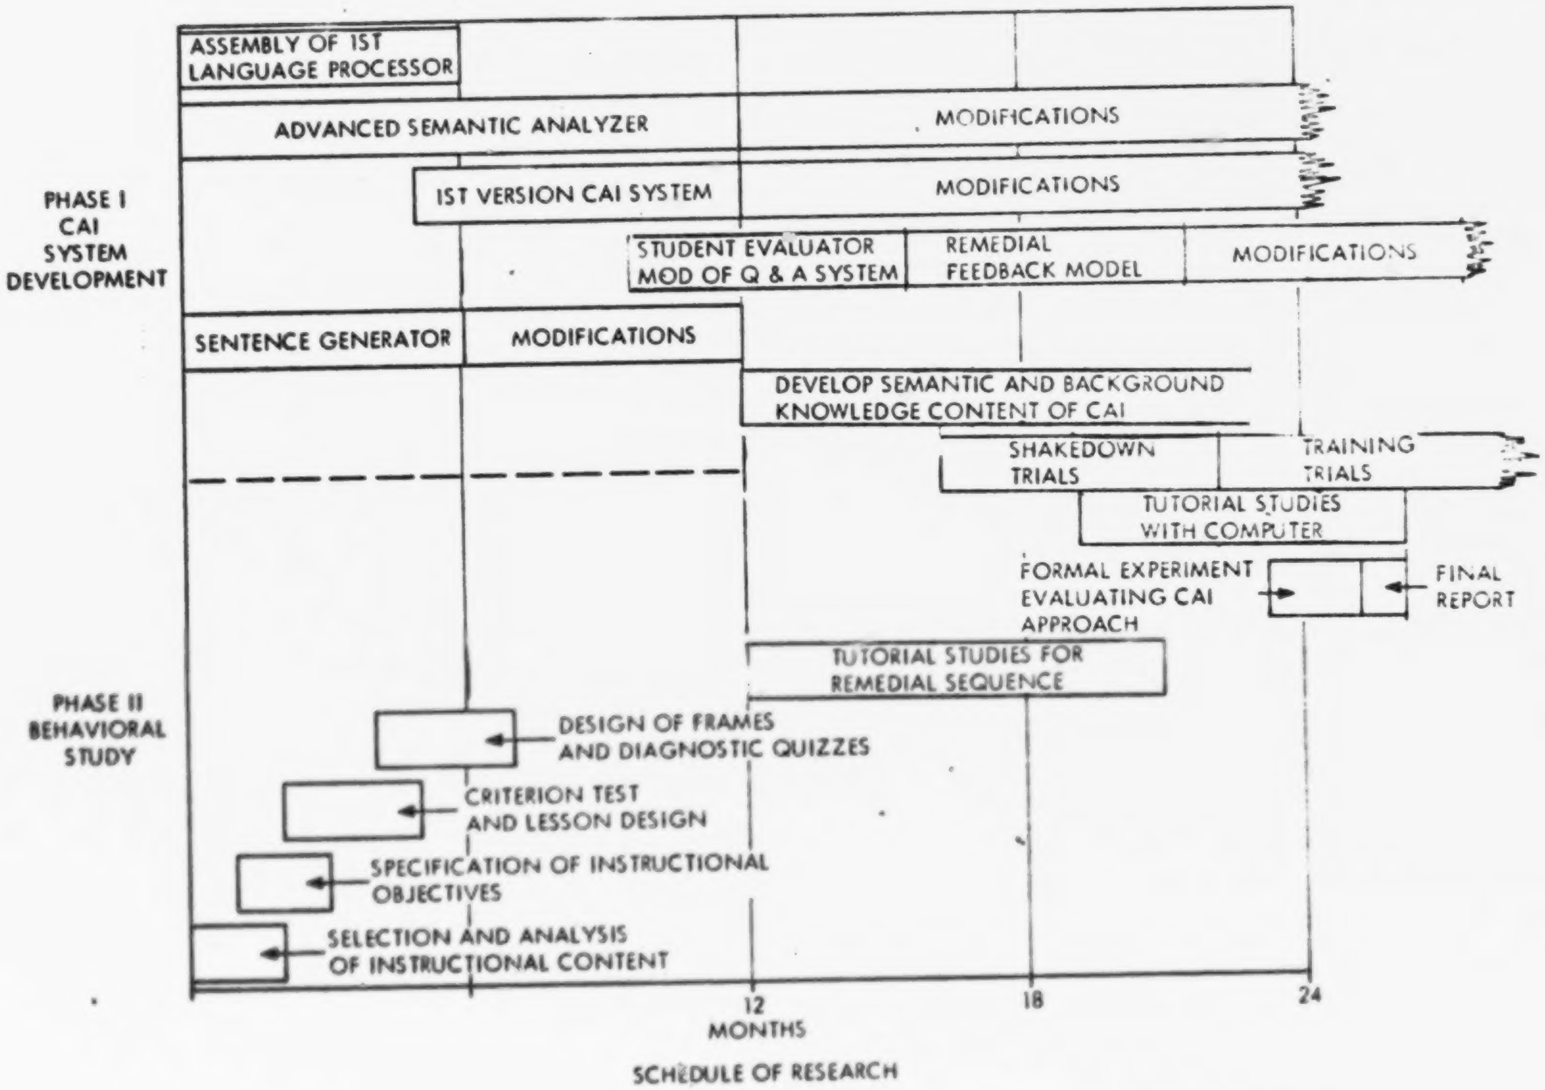

Figure 1. Schedule

Phase II

1. Selection and analysis of instructional content presumably in physiological psychology.
2. Specification of instructional objectives for a three-hour CAI lesson sequence.
3. Construction of criterion test covering the content of the entire sequence. Alternate forms of this test used as a basis for evaluating initial and terminal knowledge of students. Design of lesson strategy, chunking of text material, etc.
4. Construction of lesson frames and of diagnostic quizzes for each frame.
5. Tutorial studies simulating CAI system for purpose of developing best approaches to remedial feedback.
6. Tutorial studies using CAI system to tune it for actual use in student training. This step is largely concurrent with step 9 of Phase I and is devoted to better lesson development strategies where the Phase I operation is attempting to develop improved software for the system.
7. Conduct formal experiment evaluating effectiveness of responsive remedial approach used in CAI system.
8. Prepare final report describing CAI system and the results of experimentation in its development.

7. REFERENCES

Abelson, R. P. and Canoll, J. D. Computer simulation of individual belief systems. The American Behavioral Scientist, May 1965, 9, 24-30.

Allport, F. H. Theories of perception and a concept of structure. New York: John Wiley & Sons, 1955.

Bobrow, D. G., Fraser, J. B., and Quillian, M. R. Automated language processing. In Cuadra, C. A. (ed.) Annual Review of Information Science and Technology. New York: Interscience Publishers, 1967, in press.

Burger, J. F., Long, R. E., and Simmons, R. F. An interactive system for computing dependencies, phrase structures and kernels. SDC document SP-2454, April 28, 1966. 28 pp.

Colby, K. N. Computer simulation of change in personal belief systems. Behavioral Sciences, 1967, in press.

Coulson, J. E. Research with a program on geometric inequalities. SDC document TM-895/102/00, June 26, 1964. 84 pp.

Kellogg, C. H. (a) On-line translation of natural language questions into artificial language queries. SDC document SP-2827, in press.

Kellogg, C. H. (b) CONVERSE--A system for the on-line description and retrieval of structured data using natural language. To be presented at the IFIP/FID Conference on Mechanized Information Storage Retrieval and Dissemination, Rome, Italy, June 1967.

Klein, S. and Simmons, R. F. Syntactic dependence and the computer generation of coherent discourse. Mechanical Translation, 1963, 7(2), 50-61.

Kuno, S. Computer analysis of natural languages. Presented at Symposium on Mathematical Aspects of Computer Science, American Mathematical Society, New York, April 5-7, 1966.

Melaragno, R. J. Research with an elementary arithmetic program. SDC document TM-895/101/00, June 4, 1964. 28 pp.

Newell, A. and Simon, H. A. GPS, a program that simulates human thought. In Feigenbaum, E. & Feldman, J. (eds.) Computers and Thought. New York: McGraw Hill, 1963.

Newmark, G. Research in programmed instruction in Spanish with seventh-grade students. SDC document TM-895/103/00, June 29, 1964. 102 pp.

Olney, J. C. Some patterns observed in the contextual specialization of word senses. SDC document TM-1393/000/00, August 1, 1963. 55 pp.

Quillian, M. R. Semantic Memory. Ph.D. Thesis, Carnegie Institute of Technology, Pittsburgh, Pennsylvania.

Schwarcz, R. M. Steps toward a model of linguistic performance: A preliminary sketch. The RAND Corporation Memorandum RM-5214-PR, Santa Monica, California, January 1967.

Silberman, H. F. Exploratory research on a beginning reading program. SDC document TM-895/100/00, June 15, 1964. 190 pp.

Silberman, H. F., Coulson, J. E., Melaragno, R. J., and Newmark, G. Final report: 'Use of exploratory research and individual tutoring techniques for the development of programming methods and theory. SDC document TM-895/200/00, June 29, 1964. 58 pp.

Simmons, R. F. Automated language processing. In Cuadra, C. A. (ed.) Annual Review of Information Science and Technology. New York: Interscience, 1966, Vol. 1, 137-169.

Simmons, R. F. Answering English questions by computer: A survey. Communications of the ACM, 1965, 8(1), 53-70.

Simmons, R. F. and Londe, D. L. Namer: A pattern-recognition system for generating sentences about relations between line drawings. ACM Proceedings of the 20th National Conference, 1965, 162-175.

Simmons, R. F., Burger, J. F., and Schwarcz, R. A model of cognitive structure for verbal understanding. SDC document SP-2841, in press.

Simmons, R. F., Burger, J. F., Long, R. E. An approach toward answering English questions from text. Proceedings of the 1966 Fall Joint Computer Conference, 357-363.

Wang, H. Toward mechanical mathematics. IBM Journal of Research Development, 1960, Vol. 10, 2-22.

Weizenbaum, J. Eliza--A computer program for the study of natural language communications between man and machine. Communications of the ACM, January 1966, 36-45.

Woods, W. A., Jr. Semantic interpretation of English questions on a structured data base. Mathematical Linguistics and Automatic Translation, Report NSF-17, The Computation Laboratory, Harvard University, Cambridge, Massachusetts, 1966, 39 pp.

APPENDIX I

The attached material is Part I of a two-part paper outlining the Cognitive Structure Model for Verbal Understanding. Additional materials on semantic analysis and experimental work with the system make up the content of Part II which is not yet available in final form.

A Cognitive Structure Model for Verbal Understanding

R. F. Simmons and J. F. Burger

I. INTRODUCTION

Both in the phylogenetic and ontogenetic development of organisms, the experiencing, remembering and understanding of some aspects of their environment precede the capability to use signs and symbols to represent their experience. Most animals give obvious evidence of understanding their environment without any great capability at all for symbolic behavior. Children, long before comprehending their first words, have developed concepts of self, inside, outside, the ideas of objects, of movements, and of many relations that can hold among these concepts.

This primary capability to experience, remember, and understand--the ability to know something of the world--defines the term cognition. It is our thesis that underlying any explanation of verbal understanding there must be described a model of cognitive structure that can account for an organism's ability to perceive, recognize, and remember events and relations among the events. Once given a basic cognitive structure, the strings of natural language can be explained as a one-dimensional representation of events and relations in that structure. The idea that a natural language is a channel communicating patterns of events and relations from one such structure to another becomes a meaningful one, pregnant with the challenge of decoding linguistic patterns into the forms of the underlying cognitive structures.

II. BACKGROUND

The most recent work of structural linguists such as Chomsky [1965], Katz [1964] and others and complementary work by psycholinguists such as Miller [1965], McNeil [1966] and others has focused attention on "deep structures" underlying the obvious syntax of expressions in natural languages. The psycholinguistic work has given a strong indication that the deep structures developmentally precede ordinary language use and, in addition, are apparently closer to underlying patterns of thought (McNeil, pp. 40-62). Linguists and psycholinguists have advanced compelling arguments to show that learning and using a natural language requires far more structure than are provided by simple S-R models, Markov chains, and the like. Osgood [1963] and Miller [1965] summarize these arguments and Osgood is able to integrate an S-R probability approach at each hierarchical level of selecting components in his structural model for generating and understanding sentences.

In addition to deep structures that represent the pyramiding of simple forms into the complexity of natural language sentences, the structural linguists have also shown much concern with the content and structure of lexical entries that can be used to characterize words and other forms in a language. From a generative viewpoint the selection of certain words restricts the choice of the words that follow. For example, selecting the word "rock" as a noun-subject generally eliminates the possibility of such verbs as "see," "breathe," "eat," etc.

The linguist would like to see this kind of information associated with words and forms in the lexicon. At the semantic level, even more detailed properties are required to be associated with words to permit the selection of a particular (dictionary) sense in which a word in context is used. At both syntactic and semantic levels, linguists are now strongly of the opinion that these properties are not simple categories to which words can be assigned, but structured organizations of properties that guide their selection and interpretation (Sparck Jones [1964], Bolinger [1965], Chomsky [1965], Katz [1964]).

These lines of development all lead to the hypothesis that underlying an organism's ability to use and understand natural languages, there must exist a complex structure of information concerning properties of linguistic forms, and interrelations of words and knowledge of the world. At the semantic level Katz and Fodor [1963] offer an approach, generally seen to be unsatisfactory by Sparck Jones [1964], Bolinger [1965] and others, toward accounting for how a particular meaning is assigned to a sentence. Thompson [1966], Craig (Craig et al. [1966]) and Kellogg [1967] deal with the meaning of a sharply limited class of sentences in terms of the contents of a structured data base and introduce the idea that semantic analysis of English sentences is a process of successively mapping words and phrases into that structure. Most recently, Woods [1966] has added some generality and additional content to that type of semantic approach. Quillian [1966] in his model of human memory has taken a significant step toward showing how the meanings of words can be structured as a set of interrelations with other words that are used to define them.

Another aspect of meaning, that of inference structures, has been studied and modeled by Raphael [1964], F. Black [1964], D. Bobrow [1964] and most recently by Slagle [1965], Elliott [1965], Woods [1966] and Simmons et al. [1966]. These researchers have shown that the relational meaning of certain concepts concerning direction, part-whole, subset-superset, and numerical relations can be represented by fairly simple transformational rules of inference. The last three researches cited have shown fairly clearly how these inference structures relate to units of natural language.

These lines of linguistic, psychological, and language processing research strongly indicate that an underlying structure that would account for various kinds of understanding required in verbal comprehension must be characterized by at least the following properties:

(1) It should reflect deep relational structures that underlie the surface structure of language.

(2) It should represent meanings both in the sense of properties associated with words as required by linguists and semanticists and by the association of meanings with other related ideas.

H,

(3) It should be able to represent inference structures that allow one word or phrase to imply another, or one structure to imply another equivalent one.

A theory of verbal understanding based on such a structure should account for transforming strings of natural language into nested relational structures whose meaning is explicitly represented both as interrelationships with other structures and as related to an appropriate subset of rules or inference. Such a theory should account for important aspects of syntactic and semantic analysis of natural language. It should show how question answering, paraphrasing, and in general verbal problem solving can be accomplished. In addition it should show how meaningful and grammatical strings of language can be generated from meaning structures in the cognitive model.

In this paper we outline such a theory of verbal understanding. First we develop a model of cognitive structure that is sufficient to account for a person's ability to represent and understand the meaning of a wide range of natural language expressions. The structural theory of verbal understanding is based on this model. It includes syntactic and semantic components for transforming from English sentences into the formal language that represents the cognitive structures of the model. An explanation of question answering is presented in terms of a procedure that can accumulate inference rules for solving verbal problems or for answering questions concerning both explicit and implied relations among events. This is the central component of the theory. A system for generating natural English text from the cognitive structure model is the final component.

The model and the theory are realized in a prototype set of computer programs that accept English text and questions, transform these into formal structures of the cognitive model, use inference rules to operate on the data structure to try to answer the questions, and finally generate English statements corresponding to the data structures of any answers that may be found. The system is programmed in LISP for the AN/FS Q-32 time-shared multi-access computer system. Experiments with these programs will be described and illustrated.

III. THE COGNITIVE STRUCTURE MODEL

The elements of the model are events and relationships. The cognitive structure can be represented as a complex network whose nodes represent events and whose labelled connections or links represent relations among the events. An event-relation-event combination defines another event and a relation may itself be treated as an event. The structure is thus hierarchic and recursive. The result is that any sized unit of the structure may be treated as an event and considered in relation to other events.

An event is thus analogous to an idea, a concept, or a perception. Take, for example, the word "cougar." My idea of "cougar" is made up of visual, auditory, tactile, etc., sensations and perceptions of what I have experienced of "cougar." It also includes my emotional and motor response tendencies and my kinaesthetic perceptions of these. This idea of "cougar" is not complete in itself; it must also include changes over time in the sensations and perceptions and it must relate the concept and elements of it to whatever other aspects of the environment were perceived in spatial, temporal, emotional, or logical relations to "cougar." If "running" is one of the response tendencies I associate with "cougar," it too can be conceived of as an idea, not essentially different in its cognitive representation as an interconnected set of events and relations. Presumably, the idea "running" is represented more heavily by motor and kinaesthetic events that in the final analysis resolve to motor events and kinaesthetic perceptions of activated muscles.

Such objects as "cougar" or "running" in the cognitive structure are close-knit sets of events and relationships that ramify indefinitely throughout the structure, Despite wide ramifications, any node is an object, and it may map into the symbols of natural language. A concept like "cougar" is represented by a word; a concept like "cougars leaping from trees," while it may be a single object in the structur requires a complex string of linguistic units to map it. These meaning units are presumably morphemes and formatives as the linguist looks at language, but for th sake of simplicity we will usually deal with words.

The mapping is such that one object or event in the cognitive structure may point to several different words no one of which represents all aspects of the cognitive object. Most words will also map onto several cognitive events. Although a word may point to a cognitive event, the meaning of the word is not the event but the event's ramifications--its web of relationships to other events and its resolution into subevents that are interrelated.

The important value of the cognitive structure for understanding language is that each linguistic unit--morpheme, word, phrase, sentence, etc.--has an object as its referent. The object is always a cognitive event. With the certainty of the existence of a referent for each word, it becomes meaningful to treat linguistic units as symbols that have denotable referents. Consequently, a semantic system for a natural language--for a particular user--becomes definable as a means for resolving a many-many mapping into an unambiguous pointing from symbol to object and object to symbol. How this mapping can be resolved is discussed and exemplified in later sections.

An Abstract Nervous System: We may also consider this model of cognitive structure from the point of view of an abstract nervous system of the type mathematical biologists have explored. Here as elsewhere, we take the phenomonological view that the only knowledge an organism can develop is derived from the activities of its own neurons. This view avoids any assumptions about the nature of the "real" (outside) environment and bases the model solely on repeated patterns of stimulated neurons.

The excitation of a single neuron is taken as the most elementary of neural events. (Below this level are chemical events, molecular events, atomic events, ad inf.) If one neuron excites another, a second event occurs and a temporal relation exists between the two. If, as is usually the case, large sets of neurons are excited in different sensory modalities and include both afferent and efferent fibers, a rich basis exists for differentiating a practically infinite set of events and relations. Indeed, the problem immediately becomes one of finding commonalities rather than differences in the stimulation. Since afferent fibers pyramid upwards in a complex nervous system, there is ample opportunity to form events at successively higher levels. In consequence, what is a bewildering myriad of elementary neural events at the sensory base of the system becomes a relatively few complex events at the peak. Thus, at some level in the system of an organism that can see and hear, the simultaneous excitation of both modalities becomes, apart from all other considerations, a "seen-heard" event.

In a comparable fashion, the relation "part to whole" is an event that relates two events from the same stimulation at different levels of the ascending network of event creation. Thus a cloud is eventually perceived as part of the larger but always co-occurring stimulation of sky. Similarity is a relation in which many of the events of two different stimulations are the same. We assume that all primary logical relations such as subset, part-whole, direction, time, etc., can be derived from considering various abstracted events in the nervous system.

In contrast to afferent fibers that pyramid upwards, efferent fibers start from few nodes and ramify downwards to result, finally, in very large complex bundles of excitation to numerous motor systems. Here a cognitive event, in this case a response tendency, can trigger a whole tree of hierarchically organized response tendencies to result finally in motor behavior. Presumably, no normal motor behavior occurs in a complex organism without associated kinaesthetic stimulation that can at each ascending level create events that can be used as feedback controls on that or related behaviors.* In this fashion perceptual (viz. afferent) and motor events can exist and co-exist at all levels of a complex nervous system. Since each event can be compared as a unit to any other, events may be considered as basic units for thinking, acting, control, etc.

By introducing an appropriate theory to account for remembering useful events or (perhaps more appropriately) forgetting non-useful events, a hierarchical structure of events, recursively defined as event-relation-event, provides a remarkably satisfying framework for most forms of organic behavior.

---

*See for example Miller, Galanter and Pribram [1960].

Our intent, however, is not to derive a cognitive model from the abstract neural level, but only to show that there is a reasonable line of thought from the concepts at the cognitive level leading recursively downward to the excitation of sensory and motor nervous tissue at the neuronal level. The task of defining in detail how a nervous system can actually use elementary neural events to build complex sensations and perceptions is one to which mathematical biologists have devoted much effort over the last two decades. (See for example McCulloch [1965].)

IV. FUNCTION OF THE COGNITIVE STRUCTURE IN UNDERSTANDING LANGUAGE

One important function of a cognitive structure in an organism that uses language is to encode meanings of morphemes, words, phrases, etc., as interrelated objects in a context of other general relations that hold among events in the perceived environment. This function implies that the cognitive structure contains substructures of syntactic and semantic knowledge and rules of inference to allow for mapping language strings into the structure, mapping portions of the structure into language strings, and testing the validity (i.e., belief value) of language statements.

How such substructure may be used to accomplish these language tasks makes up the substance of a theory of verbal understanding.

Since the cognitive model requires that all information be in the form of hierarchically recursive events, where each event is defined at the next lower level as an event-relation-event structure, one problem is to show how the information contained in natural language sentences can be transformed into these structures. Generally, English sentences are complex units of meaning in which the presence of event-relation-event structures is not obvious. The syntactic categories and the sense meanings of words taken out of context are almost always ambiguous, so even though it might be shown that one underlying structure of English sentences is the event-relation-event (E-R-E) structure, there would still remain a considerable task in revealing how the contexts in which words and phrases are embedded can be used to resolve their possible syntactic and semantic ambiguities.

Later sections of this paper will show how a method of syntactic analysis can be used to transform English sentences into nested sets of E-R-E structures and a type of semantic analysis can be used to resolve many apparent ambiguities in words and phrases. At this point in the discussion we will simply assert that, with the aid of information contained in the cognitive structure, these things can be accomplished and go on to describe the structure that makes them possible.

The elements of the cognitive structure are a lexicon and a set of numbered triplets. The triplets are always in the E-R-E format where the central term is taken as a relation. A triplet may be nested as deeply as desired. For example, the following is a complex triplet:

(((E-R-E)R-E) (E-R-E) (E-R-E))

The above illustration might represent a translation from, "Large bald men eat fresh fish" into the following formal language for the structure:

(((men SIZE large) QUALITY bald) (eat SING eat) (fish QUALITY fresh))

In this illustration the uncapitalized words are events, and the capitalized terms are primitive relations. The triplet (men SIZE large) is an event. The triplet (men QUALITY bald) is another event. The middle term of the entire expression (eat SING eat) is an event in which the relational term eat is taken as an event in the singular relationship to its base form "eat." This relational-event triplet is the middle term of the expression, and it consequently relates the two complex events ((men SIZE large) QUALITY bald) and (fish QUALITY fresh).

Elements in the lexicon include the words men, large, bald, etc., as well as the primitive relations, SIZE, QUALITY, etc. A primitive relation is defined as one whose structure has certain inference properties such as reflexivity, symmetry, transitivity, etc. (about which more will be said later).

A lexical item, if a word, has associated with it a USED-IN relation to all the triplets in which that word occurred; if a primitive relation, it has defining properties associated with it. Consequently, the lexicon can be seen to be a subset of the cognitive model with the same E-R-E structure as any other elements of the model; it is a distinguishable subset because the relational terms are always either USED-IN or PROPERTY. For each word that the structure can understand there must exist a representation as a lexical item.

Figure 1a depicts the lexicon and set of triplets resulting from the example sentence; 1b shows a directed graph representation of the same information. The directed graph for a single sentence is (not surprisingly) in the form of the exact tree implied by the nested triplet structure of the sentence. As additional information is added, for example about "large bald men," nodes such as node G3 will be found to participate in other higher-level structures, with the result that G3 becomes part of a network rather than simply a tree. Adding the sentence "large bald men love food" would add nodes: G7 (love SING love), G8 (food Q IND) and the new top node G9 (G3 G7 G8).

Figure 2 shows a schematic representation of the structure represented in 1b. Although less accurate, in that it ignores the precise lexical structure of the words and phrases, the graph of Figure 2 is sufficient to use as an expository device. Further abbreviations will be introduced to ignore number, tense and case relationships except in examples where such relations are the subject of discussion.

The structure so far described is primarily a variant representation of a relational syntactic structure of the example sentences excepting only the semantic task of determining such relations as SIZE, QUAL(ity), etc. More is obviously required to model an understanding of the example sentences. If we now add information that a man is a male human; a human is an animal; to eat is to take in food; bald is a quality of lacking hair, fur, or feathers; fish is an aquatic animal; and fresh is a kind of newness and purity, the model of Figure 3 results.

Adding this new information has required the notation of such new primitive relations as SUP(erset), ASSOC(iation), EQUIV(alence) and OBJ(ect). For these to be meaningful to the model, each must be examined to determine a set of properties that may be useful in making inferences with the model. For example, let us define SUP logically as transitive, nonreflexive, asymmetric, and having

Large bald men eat fresh fish.

((((man PL men) SIZE large) QUAL bald) (eat SING eat) (fish QUAL fresh))

1. man U-I G1, G2, G3, G6
2. men U-I G1
3. large U-I G2
4. bald U-I G3
5. eat U-I G4, G6
6. fish U-I G5, G6
7. fresh U-I G5
8. PL PROP SYM
9. SING PROP SYM
10. SIZE PROP NIL
11. QUAL PROP NIL

G1 (man PL men)
G2 (G1 SIZE large)
G3 (G2 QUAL bald)
G4 (eat SING eat)
G5 (fish QUAL fresh)
G6 (G3 G4 G5)

Figure 1a. Relational Triplets

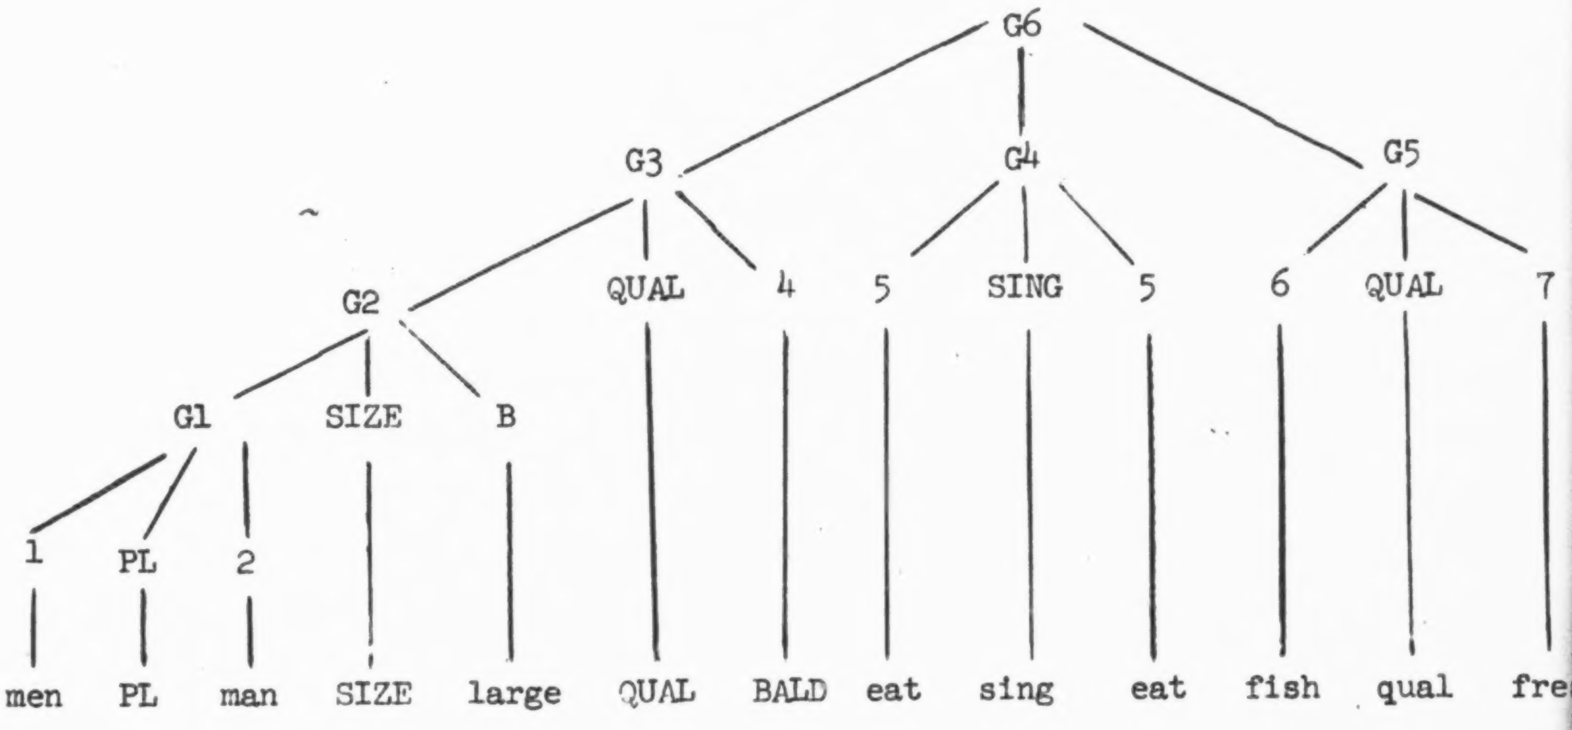
G6
G3
G4
G5
G2
QUAL
4
5
SING
5
6
QUAL
7
G1
SIZE
B
1
PL
2
men
PL
man
SIZE
large
QUAL
BALD
eat
sing
eat
fish
qual
fre

Figure 1b. Graph Structure

Relational Triplets for a Sentence

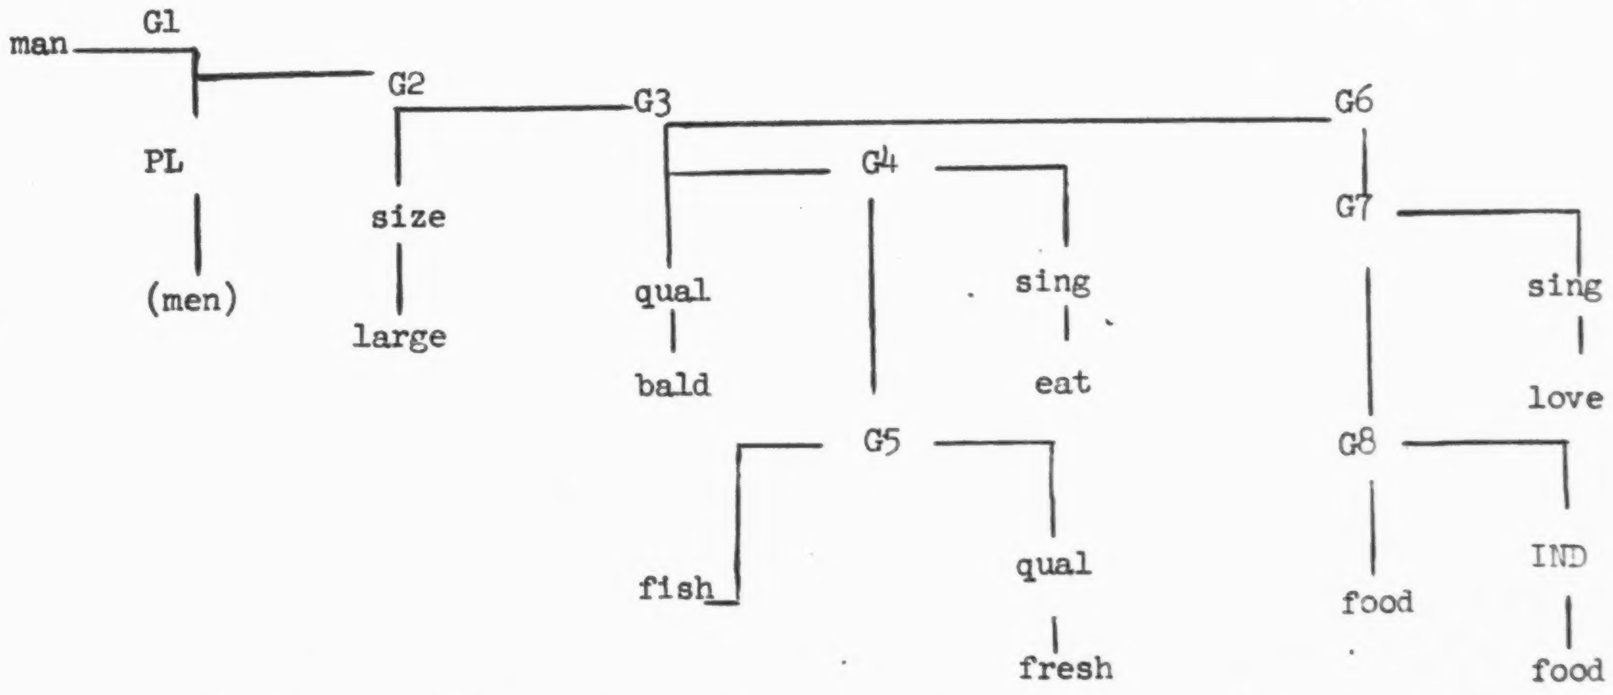
man
G1
PL
(men)
G2
size
large
G3
qual
bald
G4
sing
eat
G5
fish
qual
fresh
G6
G7
sing
love
G8
food
IND
food

Figure 2. Abbreviated Representation

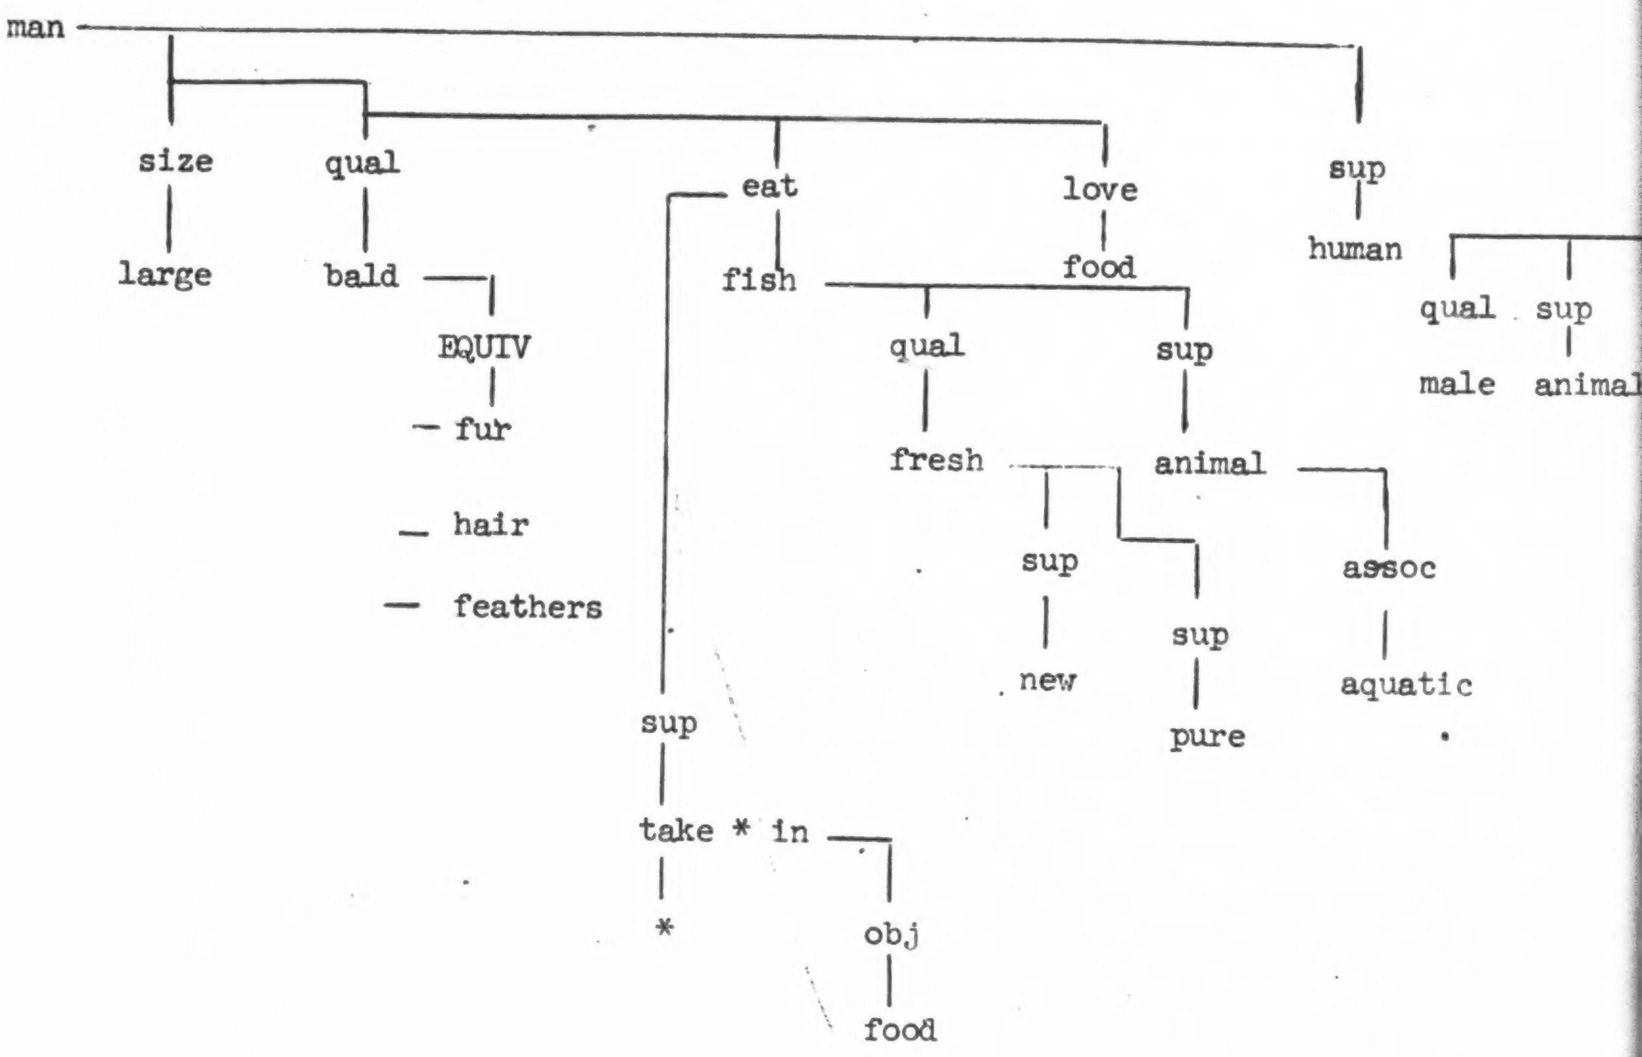
man
size
large
qual
bald
EQUIV
fur
hair
feathers
eat
fish
sup
take * in
*
obj
food
qual
fresh
sup
new
sup
pure
sup
animal
assoc
aquatic
love
food
sup
human
qual
male
sup
animal

Figure 3. Adding a Context of Knowledge

as its inverse a SUB(set), which has a similar set of properties. It is now possible to infer from Figure 3 that a hairless, furless, featherless male human animal takes in as food new, pure, aquatic animals. Significant meaning has thus been added to the statement that a bald man eats fresh fish. Exactly how this inference is carried out is discussed in a following section (p.16). For the moment, the selection and modeling of an appropriate set of primitive relations is more important.

Generally, if a word or a syntactic juxtaposition signifies something of general importance to the process of semantic analysis or inference for question answering, a primitive relational term will be noted for it. Thus concepts of temporal and spatial relations are often signified by such prepositions as "at," "in," "on," "to," "from"; these relations can be grossly summarized in context by LOC(ation) and TIME or they can be more finely represented by being shown in a SUP relation to LOC or TIME. In either case the properties of LOC or of TIME will allow certain inferences to be made that are not obvious in the meaning of the particular preposition.

Although some relations such as SIZE and QUAL may resist definition in terms of the usual logical properties of symmetry, transitivity, etc., there appear to be syntactic and semantic properties that give reason for maintaining them as system primitives. For example if a word is in a size relation to another structure, that structure can only be represented by numbers and units of measure or by a small class of size words. Qualities in general (and note later that SIZE is a kind of quality) have the syntactic-semantic property of being modified in intensity by use of certain quantifier-intensifier terms such as moderately, very, etc., or by the use of the comparative form. Eventually these may be interpretable as logical properties; at the moment they are empirically useful.

Representing Linguistic Information: In discussing the lexicon as a subset of the cognitive structure and in our representation of (men SING man) and (eat SING eats), we have hinted at the potential for representing linguistic information in the model. Linguistic data of many types can be treated in precisely the same fashion as knowledge of the environment. Figure 4 illustrates the

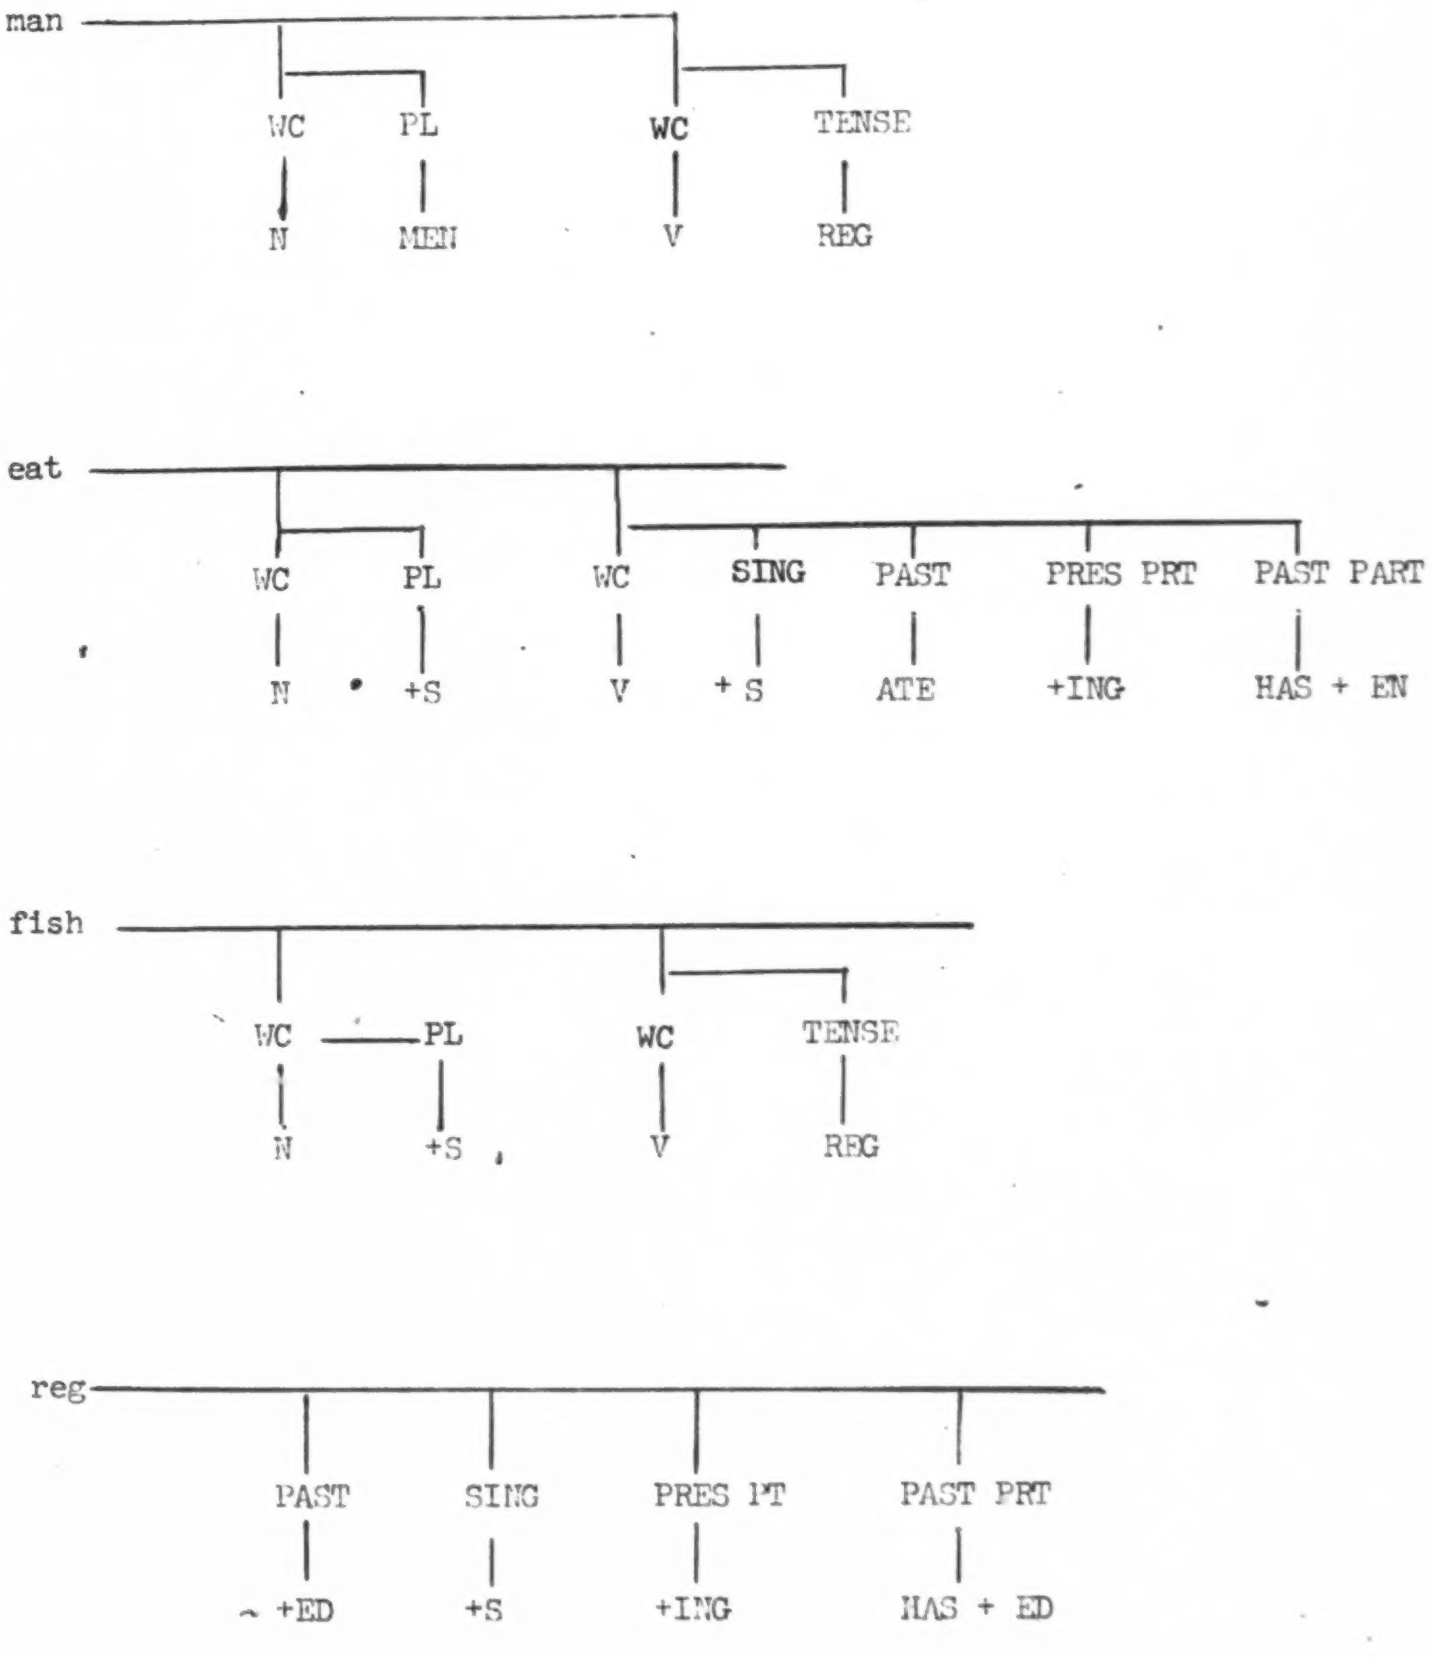
man
WC
PL
N
MEN
WC
TENSE
V
REG
eat
WC
PL
N
+S
WC
SING
PAST
PRES PRT
PAST PART
V
+ S
ATE
+ING
HAS + EN
fish
WC
PL
N
+S
WC
TENSE
V
REG
reg
PAST
SING
PRES PT
PAST PRT
+ED
+S
+ING
HAS + ED

Figure 4. Modeling Linguistic Relations

modeling of some sample linguistic information for the words "man," "fish," and "eat" taken successively as noun and verb. Under the relation "TENSE" for "man" and "fish," the term "REG" is noted. This term is defined as the English tense pattern for regular verbs. In a similar manner, +s, Has, and +ed can elsewhere be defined in a form suitable to allow the addition or stripping of "s" or "ed" as a function of preceding letters.

Such aspects of derivational morphology as rules for changing from verb to noun form by adding an -er may be encompassed by such a triple as: (V N +er). The example is not strictly true for all verbs, so it must be tied not to "V" but to some other feature related to the words for which it is true. Such syntactic-semantic notions as mass-noun, count-noun, sense-verb and the like that have found repeated usefulness in recent grammers (cf. Chomsky [1965], Katz [1964]) can be treated in a similar fashion if desired.

Quantifiers: The whole question of encoding and understanding logical quantifiers (i.e., a, an, the, each, every, none, all, some, one, two, etc.) is a thorny one. We have reached the point of recognizing that every English noun is quantified and the null article is usually to be interpreted as "generally." Whether the quantification should be associated with the noun or with the entire triplet is not yet clear to us. In either case, the means of representation would be via a QUANT relation to the particular form of quantifier. The QUANT relation would guide and limit the kinds of inference that could be performed with the triple so quantified. In the case of quantifiers, however, the coding scheme is the least of the problems: understanding the quantificational relations signified is by far more difficult. Quine [1960], Bohnert [1966] and other logicians have shed some light on the problem, but much more work is required before a full understanding is achieved.

V. VARIABLES AND RULES OF INFERENCE IN THE STRUCTURE

The complexity of inference that can be accomplished in the cognitive structure model is primarily a function of (1) how well the relational terms can be defined; and (2) the ability of the structure to represent rules of inference. The concept of variables is important both for defining relations and for representing rules of inference.* A variable symbolized as X1, X2, X3...Xn, is an object that can take as its value any other object in the structure. Thus if one wishes to express the complex logical idea of symmetry, the following formal statement can be written:

(((X1 X2 X3) AND (X2 PROP SYM)) IMP (X3 X2 X1))

This is equivalent to saying, "For any values of X1, X2, X3; if X1 is in the relation X2 to X3 and the relation X2 has the property of symmetry, then X3 is in the relation X2 to X1."

The use of variables and properties allows us to define relations with a reasonable degree of precision. A simple relation is one that can be defined by a set of properties such as transitivity reflexivity, symmetry, and others of importance to the inference system. A complex relation can either be defined as a set of simple relations or directly by a set of special rules of inference that apply to that relation.

Each of the properties used to define a relation is itself defined by an inference rule. Thus the following example definitions may be written and added like any other data to the cognitive structure:

Transitive; (((X_1 X_2 X_3) AND (X_3 X_2 X_4) AND (X2 PROP T)) IMP (X1 X2 X4))
Reflexive; (((X1 X2 X3) AND (X2 PROP REF)) IMP (X3 X2 X1))
Symmetric; (((X1 X2 X3) AND (X2 PROP SYMM)) IMP (X3 X2 X1))

Additional inference rules may be written to define other properties as they are seen to be useful. The rules can only be used in the case that the relation has the required property.

---

*We are indebted to Savitt el al. in their development of the ASP system for our understanding of the basic idea of including inference rules in the data structure. F. Black [1964], in an earlier paper, also used a variant of this idea to achieve some of the power of McCarthy's advice taker.

Relations may also be defined in terms of other relations. For example, the following are also useful in performing linguistic inferences:

((X1 (gave to X2) X3) IMP (X2 (received from X1) X3))
((X1 ((flew from X2) to X3)*) IMP (X1 ((flew to X3) from X2)*)

These inference rules in their simplest form (i.e., ((X1 fly X2) IMP (X1 (cause (move in air)) X2)) etc.) are familiar to linguists as rewrite rules; in more complex forms they are transformational rules.

There appears to be no particular limit to the number or type of variables that can be treated in complex rules and no important limitations on the generality of their application. For example, mathematical inference can be dealt with conveniently by the aid of functions such as SUM, DIFF, MULT, etc., which, in a computer representation, may be in the form of ready-made subroutines. For examples:

((FSUM X1 X2) IMP (X1 PLUS X2))
((FMULT X1 X2) IMP (X1 MULT X2))
((FCOUNT(LIST(X1 X2 X3))) IMP (How*many(X1 X2 X3)))

FCOUNT, FSUM and FMULT are to be understood as functions or subroutines that can carry out the appropriate operation and may, if desired, test to determine if the data given to them are appropriate for their operation.

Answering a question with the use of such inference rules in the data structure becomes largely a matter of trying relevant inference procedures until a successful match of the question triplet to the data triplet occurs. For example, assuming that the sentence, John kisses Mary, has been transformed into the following data structure:

1. (John kiss Mary)
2. (Kiss PROP SYMM)
3. ((X1 X2 X3) AND (X2 PROP SYMM)) IMP (X3 X2 X1))

and the question "Did Mary kiss John?" transforms to the following:

4. (Mary kiss John)

How is an answer to be obtained? First an attempt is made to match triplet 4 directly with the data structure. This fails although all elements of the query are present in the lexicon in the triplet (John kiss Mary). The relational term of the possible answer triplet is examined and its properties lead to rules of inference one of which is 3. In applying 3, Mary is substituted wherever X3 occurs, kiss wherever X2 and John for every X1. Rule 3 now reads ((John kiss Mary) AND (KISS PROP SYMM)) IMP (MARY Kiss JOHN)). The implicand is consistent with the implicator and it matches the question, so the answer is affirmative.

It can be noticed that, as a result of keying the rules of inference to named properties that are associated with particular relations, a given inference rule can only be used if it has been assigned as a property to a given relation. A more complex inference scheme such as that required for syllogistic reasoning is illustrated in the data structure of Table I.

If we assume that 1, 2, 3, and Q1 and Q2 are quantified by "all," the following procedure is used to answer the question, Q1:

a. Condor lays eggs--attempted but unsuccessful match against the data structure.

b. (condor SUP bird) AND (bird lays eggs)--discovery of a path containing all the terms of the question.

c. (SUP PROP 5)--points to inference rule #5 in data.

d. Substituting b above into rule #5, i.e., condor = X1 bird = X2, lays = X3, etc., the rule implies condor lays eggs.

e. Answer Q1 affirmative.

For Q2 the following:

a. Animal lays eggs--no match against data structure.

b. (bird SUP animal) AND (bird lays eggs)--path containing all terms of the question.

c. (SUP PROP 5)--points to rule #5.

2

Table I. A Data Structure and Two Questions

1. (bird SUP animal)
2. (bird lays eggs)
3. (condor SUP bird)
4. (SUP PROP 5)
5. (((X1 SUP X2) AND (X2 X3 X4)) IMP (X1 X3 X4))
6. (SUP INV SUB)
7. (INV PROP 8)
8. (((X1 X2 X3) AND (X2 INV X4)) IMP (X3 X4 X1))
9. (SUB PROP 10)
10. (((X1 SUB X2) AND (X2 X3 X4)) IMP ((X1 X3 X4) QUANT SOME))

Q1. (condor lays eggs)

Q2. (animal lays eggs)

d. Substitute as in Q1 but this time the rule does not match the data path.

e. No other SUP properties are indicated so take the inverse relation #6 (SUP INV SUB).

f. Inverse points to property 8, which transforms (bird SUP animal) to (animal SUB bird).

g. SUB points to PROPerty #10.

h. Substituting in #10(((animal SUB bird) AND (bird lays eggs) IMPlies ((animal lays eggs) Quantified Some))

i. Answer Q1, "Some animals lay eggs."

In Table I and the preceding explanatory use of it, linguistic-logical relations are patterned after the corresponding logical relations of set theory. If a bird is a kind of an animal, then (bird SUP animal) and (animal SUB bird) can represent this fact formally in the model. Many relations such as SUP and SUB have clearly defined inverses and the use of the inverse is one of the primary forms of linguistic inference for use in question-answering. We can also see from Table I that not only may simple relations point via properties to inference rules, but also they may exist as events in relation to other relations as in (SUB INV SUP), consequently implying the use of lower-level inference rules pertaining to that relation (i.e., INV with the property, rule #8).

Answering Q2 involved first the use and rejection of an inappropriate inference rule, then a transformation of the data by discovering an inverse relation and an inference rule associated with it and finally the use of an inference rule associated with the already transformed data. It can be seen that some complex questions might possibly require many rules of inference and many transformations on the data before an answer is discovered. One immediate problem that arises is that in discovering that no answer exists in the system, all relevant transformations and rules of inference must be tried with reference to all paths that contain elements of the question. Another problem is that of ordering the use of transformations and rules of inference. What these problems imply is that answering complex questions by using rules of inference on stored data can easily achieve or surpass the magnitude of effort required to solve a chess problem or to prove a complicated logical theorem. The tree of possible solution paths for a complex question is finite* but often very large.

Question answering from this point of view becomes a process almost identical to that followed by Newell, Shaw, and Simon [1963] in their approach to GPS, the General Problem Solver. Since the cognitive model can incorporate subroutines and functions as parts of inference rules, it like GPS can be used for solving any problems that can be translated into a structure of binary relations. Actually solving such problems requires not only the development of appropriate rules for inference, but also the discovery of tree-pruning and other heuristics to reduce the possibility space in which to search for an answer.

Thus, as in GPS, it is only theoretically true that given a sufficient data base and an adequate set of rules a pertinent question can be answered by the cognitive structure. The data may be present but the tree of possible transformations and inferences may be so large that it cannot be explored by any practical system (including organisms) in any reasonable length of time. It may be that completely parallel computing systems such as those envisaged by Savitt and his associates [1966] may so vastly shorten the time required to explore large sets of inference paths that computers might come to solve some complex problems more rapidly than people. For today, however, much evidence exists that the serial computer is intrinsically far less efficient than humans are for determining a desirable course of action from a large tree of possibilities (see Dreyfus [1965]).

---

*Except for certain kinds of recursive inference rules that can be controlled by limiting the depth of recursion that is allowed.

The fact that question-answering in this model of cognitive structure reduces to a generalized problem-solving task is encouraging support for the validity of the model. In many previous attempts at question-answering it appeared that there were hundreds of types of questions (list, name, count, who, what, when, etc.), each possibly requiring a special function to examine data for an answer. The present analysis shows that if a question can be reduced to terms of the model, one generalized procedure--essentially the same one required for any kind of problem solving--is sufficient (at least theoretically) to determine an answer. It is further encouraging, and not entirely unexpected,* that the process of question-answering and verbal understanding intersects with other problems studied by researchers in artificial intelligence and heuristic programming, namely game playing, problem solving, and theorem proving. The differences between verbal understanding and such other tasks lie mainly in the kinds of inference required to transform strings of language into nested event-relation-event structures of the cognitive model. Our approach to this problem is described in the following two sections.

---

*For example F. Black [1964] developed a general inference system as a question answerer, then realized that it could deal successfully with McCarthy's [1959] Advice Taker problem. Also Slagle's [1965] DEDUCOM (Deductive Communicator) used a similar inference system to solve Advice Taker problems, answer questions, and solve certain GPS tasks.

VI. FROM ENGLISH TO RELATIONAL KERNELS

In previous sections we have developed a model of cognitive structure and indicated its power in answering questions posed to it in a formal language equivalent to the structure. It is now necessary to show how English statements can be translated into that formal structure. In general, the theory of verbal understanding posits that the problem of understanding a natural language expression is one of discovering how the string of natural language symbols can be mapped into the formal structures of the cognitive model.

The semantics of a language is defined generally as the mapping of symbols onto the objects that they denote. Discussions of denotation are often confused by the observation that many words of a natural language do not map onto objects. For example, function words such as "the," "and," "of," and such words as "concept," "collection," etc., have no real world denotation. The function words signal various relations among other words and the abstract words are agreed-upon symbols of complex concepts. In our view, at the simplest level, every natural language word and phrase does, in fact, denote an object. The objects denoted are cognitive objects. As described earlier, a cognitive object is a node in the cognitive structure. This node may represent a simple concept as in "table" or it may reflect a tremendous range of information as in "meson" or "quasar." In fact, even the simplest concepts ramify throughout the cognitive system and thus develop essentially an open-ended richness of meaning. (See Quillian [1966] for a discussion of this point.)

The meaning of a word is thus the set of events that ramify from the node or object onto which it maps in the cognitive structure.

In English a word out of context can map onto several or many cognitive objects. Similarly, a cognitive object may be equivalently expressed by many different words or phrases. The problem of translating a string of English into the cognitive structure, or conversely, expressing an idea in English, is thus one of resolving a many-many mapping in both directions. In linguistics the problem is familiar in attempting to discover the intuitively best syntactic analysis of a sentence. In semantics the problem has been expressed by Katz and Fodor [1963] and others under the term "disambiguation."

If each word can map (on the average) onto n objects, and an English text string is m words in length, it is theoretically possible to have mn possible interpretations of the string. In fact, humans do much better than this in finding usually one (or in the case of puns, two) prominent interpretations for a given sentence. Numerous experiments (e.g., Miller [1965]) have shown that they accomplish this vast reduction of interpretation space by the use of associated contexts--verbal and perceptual, explicit and implicit.

In our approach, we assume that a listener--largely in sequential fashion--reduces a string of perceived language symbols into a nested structure of relational triplets of the same form that we posit as cognitive structures. We believe this is accomplished as one complex process that combines at each step linguistic, semantic, logical and experiential analysis. In our model, however, we still separate out a phase of syntactic processing to produce a nested set of English kernel structures followed by a semantic processing that transforms the English kernels into unambiguous relational triplets that map onto the cognitive structure. In a later section (Part II) the combination of these two phases of analysis into a single one is discussed.

VII. SYNTACTIC ANALYSIS

The role of syntactic analysis in the present model is to reduce a complex sentence such as the following:

> "The condor of North America called the California Condor is the largest land bird on the continent,"

into a set of simple nested kernels such as those below:

> ((((condor art the) of (America N* North)) called ((Condor N California) art the)) is ((((bird N land) Adj largest) art the) on (continent art the)))).

The nesting structure of these linguistic kernels is precisely the form and nesting structure of relational triplets in the cognitive structure. An English kernel, in our view, is always made up of an object word, a relational word, and an object word. The middle terms "art," "adj," and "N" are signals to the semantic system to select certain relations. The third term is frequently null, as in the case of intransitive verbs, i.e., (birds fly ***).

Our present procedure for analyzing a sentence into its syntactic kernels involves first a dependency analysis, then a conversion from the dependency structure to an immediate constituent (IC) tree structure, and finally, the use of both dependency and IC information to reduce the structure to nested kernels. Although we believe simpler approaches are possible (and desirable), the approach we use was developed prior to our model of cognitive structure. It is briefly described below. A more complete description is available in Burger, Long, and Simmons [1966].

The dependency analysis procedure requires word-class information (i.e., noun, verb, preposition, article, adjective, etc.) stored in a special dictionary.* It also depends heavily on context rules also available in the dictionary. Given a sentence such as

1. The man for whom the bell tolls is dead.

the first step is to look up each word in the dictionary to discover its word-class and context possibilities. The following set might result:

the: * ART N N , RPRON ART N N
man: ART N PREP V
for: N PREP RPRON V
whom: PREP RPRON ART *PREP

bell: ART N V V
tolls: N V V *N
is: V V ADJ *
dead: V ADJ * *V

Although the dictionary lookup would usually result in several frames for each word, only one or two are shown in this first example to help clarify the procedure for analysis. The 4-tuples associated with each word, W, show for each known context: first, the class of the preceding word; next the class of the word, W, itself; then the class of the following word; and fourth, the class of the word that can govern W in that context. In this example, "bell" is preceded by an ART(icle), is itself a N(oun), is followed by a V(erb), and can be governed by a V(erb) following it. (Being governed by a preceding V or N would be signified *V *N respectively.).

---

*It was noted earlier on p. 15 how this information can be coded into the cognitive model.

By fitting the 4-tuples together in sequence as illustrated below:

Context								Dependency
*	ART	N						N
	ART	N	PREP					V
		N	PREP	R PRON				V
			PREP	R PRON	ART			PREP
				R PRON	ART	N		N
				etc.				
					ART	N	V	V

we can accept or reject word-class possibilities on the basis of the context of the sentence being examined. As a result of this analysis, several strings representing possible word-class sequences result. These are used as candidates for making dependency analyses.

The dependency analysis is accomplished with the aid of a pushdown storage list and some tests for well-formedness. The processing is done sequentially. For example, the first word of Example 1, "the," is looking for a N to govern it. "The" is placed on a pushdown list and the next word is examined to discover if it is an N. "Man," the second word, is an N and the dependency link "the" governed by "Man" is produced; Man is put on the pushdown stack and the next word is examined to see if it is the V for which "Man" is looking. It is not, so "for" looking for a V is put on the list and so on. As each word is considered, a check is made on the word topmost on the list and the word immediately following in the sentence string. As a word finds its governor, it is popped off the pushdown list and the next word down becomes the top. Normally, one pass through the sentence is sufficient to complete the set of dependency links for all the words. Inconsistencies can result and these cause additional tests to be made. When all strings for a given sentence have been passed through the analyzer, none, one, or several dependency structures may have resulted. For the sentence, "Time flies like an arrow," the following three structures were obtained.

(1 TIME N FLIES 2)	(1 TIME V * 0)	(1 TIME ADJ FLIES 2)
(2 FLIES V * 0)	(2 FLIES N TIME 1)	(2 FLIES N LIKE 3)
(3 LIKE PREP FLIES 2)	(3 LIKE PREP TIME 1)	(3 LIKE V * 0)
(4 AN ART ARROW 5)	(4 AN ART ARROW 5)	(4 AN ART ARROW 5)
(5 ARROW N LIKE 3)	(5 ARROW N LIKE 3)	(5 ARROW N LIKE 3)
Parsing 1	Parsing 2	Parsing 3

Each element reads (sequence number of word, word, word-class, governing word, sequence number of governing word). The equivalent representation as dependency trees is shown by Figure 5 below.

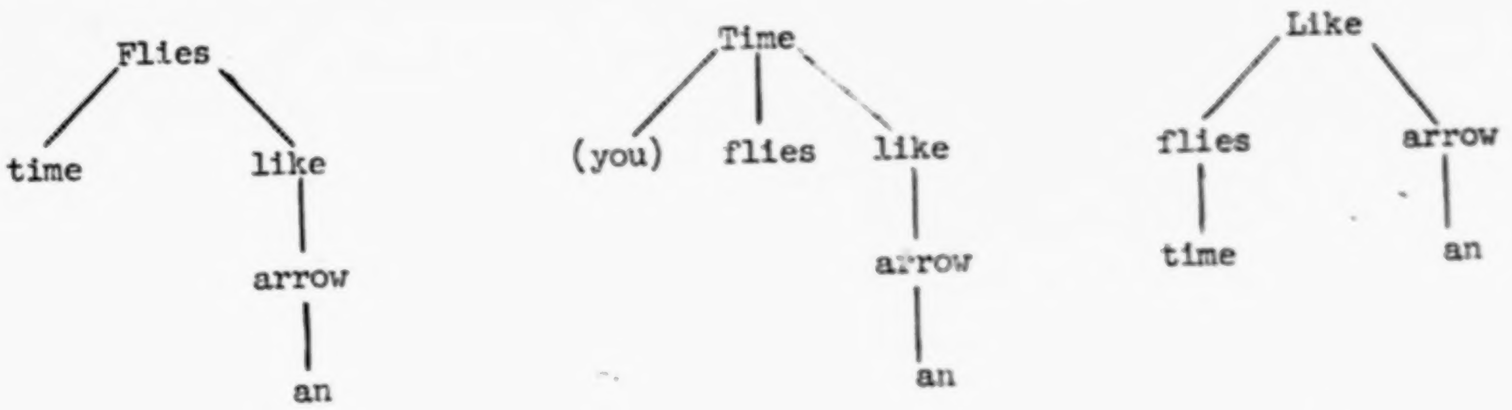

1. 2. 3.

Figure 5. Dependency Trees for "Time flies like an arrow."

VII. FROM DEPENDENCY TO IMMEDIATE-CONSTITUENT ANALYSIS

The next stage of our analysis requires that an Immediate-Constituent (IC) structure be generated. Garvin [1965] points out that IC analysis, as it is used in a recognition grammar, separates in an orderly way the set of words in a sentence into progressively smaller subsets until the final subset contains one word S for sentence. At each step, separation is made according to rules that restrict the ways in which the set, or a subset, may be divided and that provide labels for the newly formed subsets. These labels are standard linguistic terms such as noun-phrase, verb-phrase, subordinate clause, prepositional phrase, etc., that describe the use of each labelled subset as a syntactic substructure of the sentence. Any subset is then called an immediate-constituent, and the set of all labelled immediate-constituents is called an IC structure, which is a form of a phrase-structure analysis of the sentence.

An alternative, and more common, method of construction begins with the set of words in the sentence and progressively combines pairs of elements (initially the word-classes of words) to make a single element. This approach forms the basis for a computable algorithm. Analogous to the first method above, rules are applied at each step to determine which two elements shall be combined, and to apply a label to each newly formed element. Combination continues in this manner until the set consists of a single element representing the entire sentence.

IC Algorithms: The particular IC structure that we generated is based on the tree reflected by the dependency analysis and on the word-classes assigned to each word. A set of rules has been devised and is contained in an IC Rules table. While this table is too large to be shown here in its entirety, it is exemplified by the small sample shown in Figure 6.

ART + N : NP
ADJ + N : NP
ART + NP : NP
PREP + NP : PP
NP + PP : NP
V + ADV : VP
AUXV + V : VP
VP + NP : VP
NP + VP : S

NP = noun-phrase
PP = prepositional phrase
VP = verb phrase
S = sentence

Figure 6. Examples from the IC Rules Table

In combining elements for an IC analysis three conditions must be met:

(1) One of the pair of elements must be dependent on the other.
(2) The two elements must be adjacent relative to the ordering of the original sentence.
(3) There must exist a rule in the IC Rules table to define and describe their combination.

If these three requirements are satisfied, the two elements are combined and labelled with the phrase name provided by the rule. The new element then replaces this pair in the sentence string and assumes the dependency and governor relationships formerly held by the governing member of the pair.

Processing begins at the lowest dependency level, combining words at that level with their governors at the next higher level whenever the three requirements are met. Not until all words at a given level are joined with their governors does the procedure "move up" a dependency level to continue analysis.

When all possible combinations have been made at the zeroth level (e.g., the top of the dependency tree) the results are examined. For many sentences and parsings, the set will now consist of the single element (labelled "S") representing the entire sentence. If this is the case, the analysis is complete.

In many other English sentences, however, the word order is such that, if the adjacency requirement is strictly invoked, certain words or phrases may never be combined with their governors. This situation, and the way in which we handle it, is best described by an example. Consider the sentence, "When summer came, Bill painted his boat." At a particular stage of IC generation, the introductory phrase "When summer came" will have been combined and labelled as SC (subordinate clause), "Bill" will still be a noun, "painted" a verb, and "his boat" will have been combined into an NP (noun-phrase). The verb is dependent at the zeroth level (the top) and all other words and phrases at this point are dependent on it. Now if the adjacency requirement is continually enforced, the verb and NP will combine to make the VP (verb-phrase) "painted his boat" followed by the combination N + VP = S to create the element "Bill painted his boat" labelled "S." The SC still precedes this element and, while it is now found to be dependent on it and the two are adjacent, there is no IC rule to depict the combination "SC + S." The two cannot be combined.

Recognizing the need to deal with these "isolated ICs" at this point, we override the adjacency requirement by applying transformations to reorder the partially completed IC structure. In the example cited, the transformation rule applied would move the SC between the verb and the NP, thus reordering the sentence to read, "Bill painted after summer came his boat." While, as a spoken English sentence, this ordering is awkward, the IC procedure can now reduce the structure to a single element that seems properly to describe the phrase structure of the sentence.

The extent to which transformational rules are required is not yet wholly clear to us. The optimal format for these rules is also still indeterminate. It is clear that any dependency structure in which the constituents of phrases are separated by one or more words requires some form of transformational operation to make a continuous phrase structure tree. The transformations can be applied literally to the ordering of the words of the dependency-analyzed sentence as illustrated above, or they can be applied at some higher level in the tree as is often done in ordinary transformational analysis approaches used by Zwicky, et al. [1965] or Petrick [1965]. Further research will clarify this problem.

Output of the IC analysis program is presented in the parenthetical notation of LISP (see Figure 7), or on a display scope as a labelled tree structure (see Figure 8).

```
(S (NP (ART THE)
       (NP (N BOOK)
           (S (RPRON THAT) (S (PRON YOU) (V READ)))))
       (VP (V IS)
       (PP (PREP ON)
           (NP (ART THE)
               (NP (TABLE)
                   (PP (PREP IN) (NP (ART THE) (N HALL))))))))
```

Figure 7. Nested Representation of IC Analysis for the Sentence, "The book that you read is on the table in the hall."

Transforming from the IC structure into nested kernels is a relatively simple process of looking up each IC triplet (i.e., NP + Art, N; S = NP, VP, etc.) to discover if it transforms into a kernel structure. Thus NP = Art, N transforms to (N art Art) and NP + N, N transforms to (N n N). The lower case symbols "art" and "n" are relational terms to be passed onto the semantic analysis system. The upper case represent the word to which the word-class was assigned. If we consider the IC analysis of example Sentence 1 on page 25, the following rules are sufficient to generate the nested set of relationship kernels.

NP = N_1, N_2 - N_2 n N_1

PP = Prep, NP - 0

NP = N, PP - N PREP NP

NP = Art, NP - N art ART

VP = V, NP - 0

S = NP, VP - N V (NP)

S = S, VP - N V (NP)

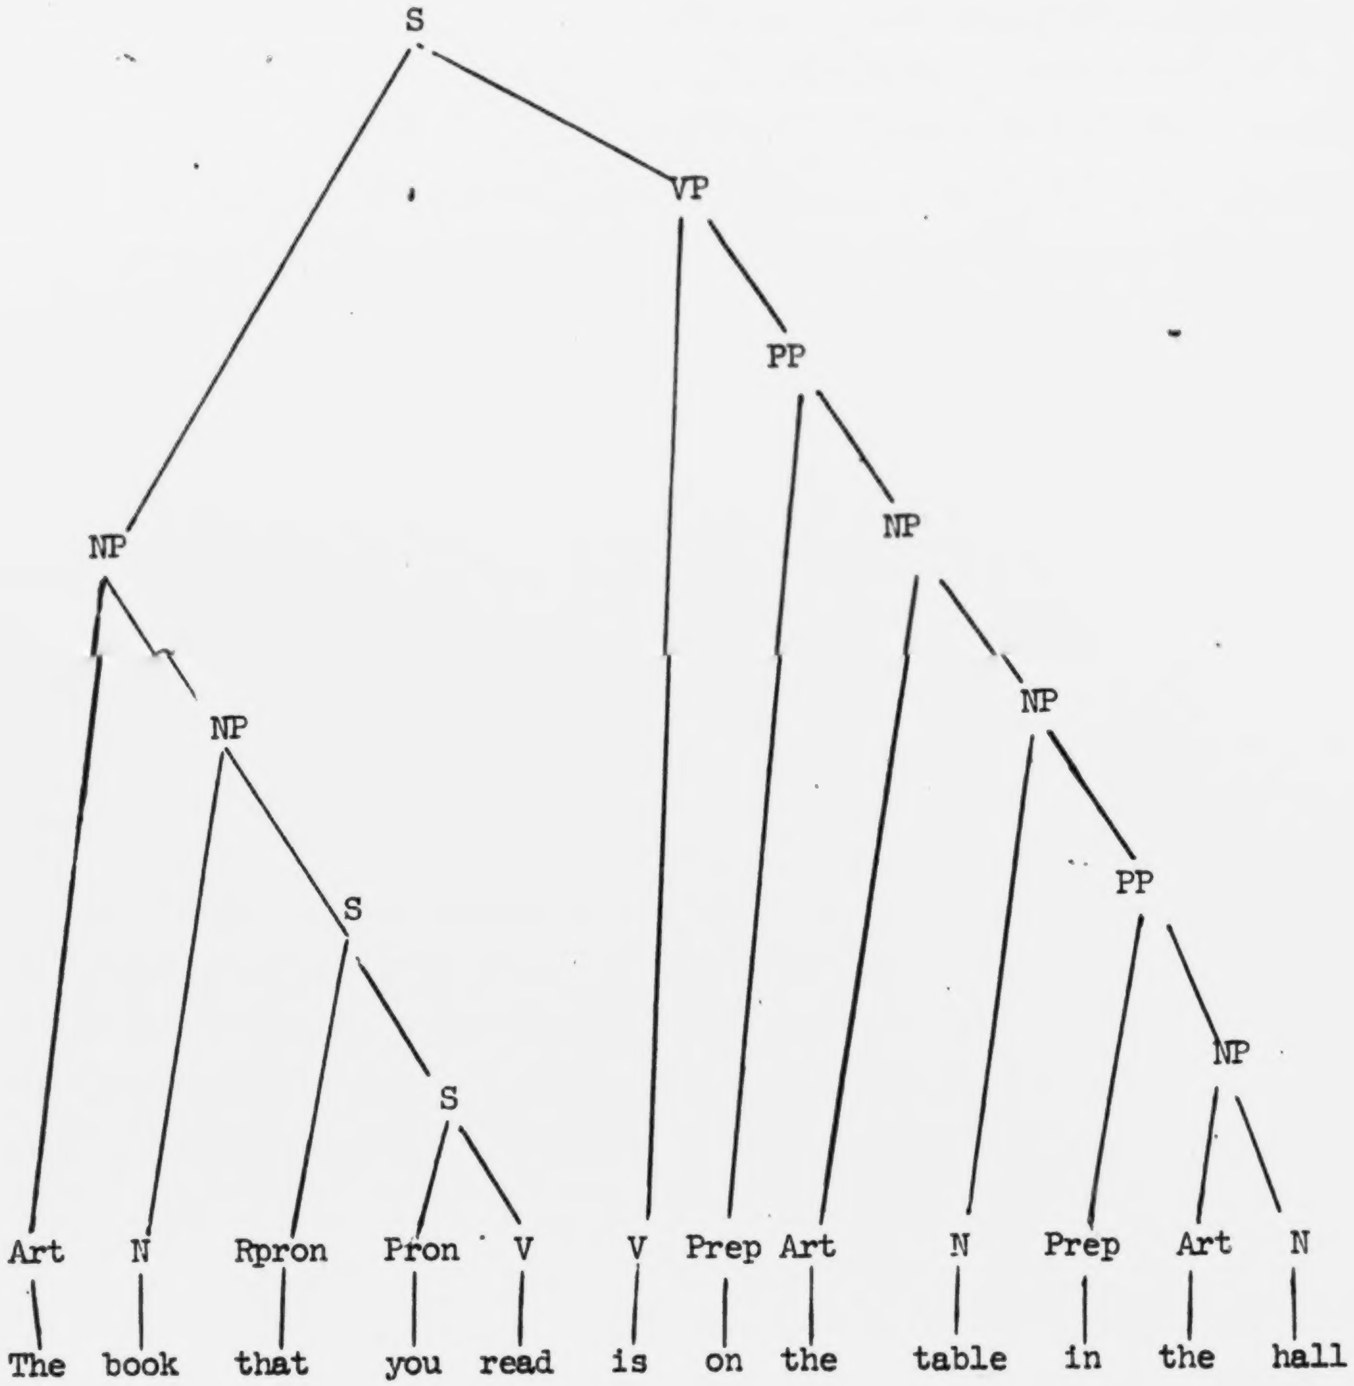
S
VP
PP
NP
NP
NP
NP
S
PP
S
NP
Art
N
Rpron
Pron
V
V
Prep
Art
N
Prep
Art
N
The
book
that
you
read
is
on
the
table
in
the
hall

Figure 8. Display Output of IC Generator for the Sentence, "The book that you read is on the table in the hall."

Parentheses surrounding a kernel term indicate that it is optional, depending in these cases on whether or not the verb phrase contains an NP. The nesting is obtained by respecting the nested structure of the IC analysis using the kernels as lowest level units.

More complex rules are required for deriving the kernels from conjunctive and infinitive constructions but in all cases they are relatively simple transformational rules. The kernels that result from the example sentence are repeated below:

> "The condor of North America called the California Condor is the largest land bird on the continent,"
>
> ((((condor art the) of (America N* North)) called ((Condor N California) art the)) is ((((bird N land) Adj largest) art the) on (continent art the))).

Sentences of great variety have been used as experimental inputs to this system. The performance is generally rapid and the output quite satisfactory for additional processing in the language model. It is quite obvious to us that the PLP-II syntactic analyzer is far more complex than the system required merely to furnish bracketings of nested structures of English sentences, but rather than patch, simplify or rewrite PLP-II, we prefer to devote our efforts to developing the semantic analyzer presented in Part II. It is our expectation that our semantic system will eventually encompass the syntactic approach to result in a single operation that transforms from English strings directly into unambiguous formal structures of the model.

REFERENCES

Black, F. S. A deductive question-answering system. Ph.D. Thesis, Div. Eng. Appl. Phys., Harvard University, Cambridge, Massachusetts, June 1964.

Bobrow, D. G. Natural language input for a computer problem-solving system. Proceedings of the Fall Joint Computer Conference, 25, 1964. Also available as Ph.D. Thesis, Math. Dept., M.I.T., Cambridge, Massachusetts, 1964.

Bohnert, H. G., and Becker, P. O. Automatic English-to-logic translation in a simplified model. I.B.M., Thomas J. Watson Research Center, Yorktown Heights, New York, March 1966. AFOSR 66-1727 (AD-637 227). 117 pp.

Bolinger, D. The automization of meaning. Language, 41, October-December 1965. 555-573.

Chomsky, N. Aspects of the Theory of Syntax. Cambridge, Massachusetts: M.I.T., 1965, 251 pp.

Chomsky, N. Syntactic Structures. The Hague: Mouton, 1957, 116 pp.

Craig, J. A. DEACON: Direct English access and control. Proceedings of the Fall Joint Computer Conference, Washington, D.C.: Spartan, 1966, 365-380.

Dreyfus, H. L. Alchemy and artificial intelligence. The RAND Corporation, P-3244, Santa Monica, California, December 1965, 90 pp.

Elliott, R. W. A Model for a Fact Retrieval System. Ph.D. Thesis, University of Texas Computation Center, May 1965, TNN-42.

Katz, J. J. Analyticity and contradiction in natural language. In: J. A. Fodor and J. J. Katz (eds.) The Structure of Language: Readings in the Philosophy of Language, Englewood Cliffs, New Jersey: Prentice-Hall, 1964.

Katz, J. J., and Fodor, J. A. The structure of a semantic theory. Language, 1963, 39, 170-210. Reprinted in part in: J. A. Fodor and J. J. Katz. The Structure of Language: Readings in the Philosophy of Language. Englewood Cliffs, New Jersey: Prentice Hall, 1964.

Kellogg, C. H. On-line translation of natural language questions into artificial language queries. SDC document SP-2827, in press.

Kellogg, C. H. CONVERSE--A system for the on-line description and retrieval of structured data using natural language. To be presented at the IFIP/FID Conference on Mechanized Information Storage Retrieval and Dissemination, Rome, Italy, June 1967.

McCarthy, J. Programs with common sense. Proceedings of the Symposium on Mechanization of Thought Processes, I, London, England: H.M. Stationary Office, 1959, 75-91.

McCulloch, W. S. Embodiments of Mind. Cambridge, Massachusetts: M.I.T., 1965.

McNeill, D. Developmental psycholinguistics. In: F. Smith and G. A. Miller (eds.) The Genesis of Language. Cambridge, Massachusetts: M.I.T., 1966.

Miller, G. A. Some preliminaries to psycholinguistics. American Psychologist, 1965, 20, 15-20.

Miller, G. A., Galanter, E., and Pribram, K. H. Plans and the Structure of Behavior. New York: Holt, 1960.

Osgood, C. E. On understanding and creating sentences. American Psychologist, 1963, 18, 735-751.

Petrick, S. R. A Recognition Procedure for Transformational Grammars. Ph.D. Thesis, M.I.T., Department of Modern Languages, Cambridge, Massachusetts, June 1965.

Quillian, M. R. Semantic Memory. Ph.D. Thesis, Carnegie Institute of Technology, Pittsburgh, Pennsylvania, February 1966.

Savitt, D. A., Love, H. H., Troop, R. E. Association-storing processor study. Technical Report RADC-TR-66-174, Rome Air Development Center, Rome, New York, June 1966.

Schwarcz, R. M. Steps toward a model of linguistic performance: A preliminary sketch. The RAND Corporation Memorandum RM-5214-PR, Santa Monica, California, January 1967.

Simmons, R. F., Burger, J. F., Long, R. E. An approach toward answering English questions from text. Proceedings of the 1966 Fall Joint Computer Conference, 357-363.

Slagle, J. R. Experiments with a deductive question-answering program. Communications of the ACM, December 1965, 8, 792-798.

Sparck Jones, K. Synonymy and Semantic Classification. Cambridge Language Research Unit, Cambridge, England, June 1964, M.L. 170.

Sparck Jones, K. Semantic Markers. Cambridge Language Research Unit, Cambridge, England, November 1965, M.L. 181. (Mimeographed.)

Thompson, F. B. English for the computer. Proceedings of the Fall Joint Computer Conference, Washington, D.C.: Spartan, 1966, 349-356.

Woods, W. A., Jr. Semantic interpretation of English questions on a structured data base. Mathematical Linguistics and Automatic Translation, Report NSF-17, The Computation Laboratory, Harvard University, Cambridge, Massachusetts, 1966, 39 pp.

Zwicky, A. M., et al. The MITRE syntactic analysis procedure for transformational grammars. AFIPS Conference Proceedings, Vol. 27, Part 1, 1965 Fall Joint Computer Conference, Washington, D.C.: Spartan, 1965, 317-326.

ii

APPENDIX II

Sample of Minimal Lesson Structure*

FRAME 1

The primary receptor cells of the retina in man are of two discrete types: the cones, concentrated mostly in the centre in the fovea, and the rods located outside this area. The greater the distance from the fovea the smaller the ratio of cones to rods, until in the extreme peripheral field scarcely any cones are found. The names derive from the microscopic appearance of the two types of cell and are more aptly descriptive of the shapes found in some animal eyes than in the human eye, but the principal functions of the two types are more distinct.

a. Student should know the names of the two types of retinal receptor cells.

b. Student should show understanding of the general areal distribution of the rods and cones relative to the fovea.

FRAME 2

The cones are only slightly responsive to changes in intensity of light, and in fact need considerable threshold intensity before they will react at all, but they are extremely sensitive to outline and to movement; they are also the principal receptors for colour vision in man and in those animals which are not colour-blind. There is no very good evidence that the common laboratory animals, such as rabbits, cats, and dogs, have colour vision, nor, in spite of all the tales told by the afficionado, has the bull. Primates are, in fact, the only mammals other than man in whom colour vision has been definitely proved, although it has been demonstrated beyond reasonable doubt in several insects, fishes, and birds.

---

* The text in this sample lesson is taken with modification from The Electrical Activity of the Nervous System, Mary A. R. Brazier, Macmillan, New York, 1953.

a. Student should understand the function of the cones with respect to:

reaction to changes in light intensity;
reaction to absolute light intensity;
sensitivity to outline and movement;
colour vision.

FRAME 3

The rods serve a different purpose from the cones and react maximally to a different stimulus: they are very sensitive to light, having a low threshold for intensity of illumination and reacting rapidly to a dim light or to any fluctuation in the intensity of the light falling on the eye. This differentiation of two types of end-organ in the eye, each with a distinct function, is the essence of the duplicity theory of vision as originally formulated by Schultze and later by von Kries.

a. Student should understand difference between function of rods and cones with respect to:
intensity of light (threshold);
changes in light intensity;
colour vision.

b. Student should be able to say whether the fovea or the peripheral field is more sensitive to changes in light intensity.

FRAME 4

The innervation of these two types of end-organ is also different structurally. In the centre of the human fovea, where there are no rods, the cones are each innervated through a bipolar neuron by the sole dendrite of a ganglion cell whose axon runs directly in the optic nerve to the optic thalmus; they thus convey exactness of detail. In reptiles and birds, especially hawks, which have great visual acuity, the fovea is very highly developed. By contrast with the cones, the rods do not have individual innervation, for several of these receptors are found to connect with multiple dendrites of a common bipolar cell. In the extreme peripheral field as many as 200 rods may make synaptic connection with a single bipolar cell. Thus the impulse reaching a ganglion cell from a cone in the fovea is from an exactly circumscribed area of the retina and conveys detailed information, whereas an impulse in a nerve cell whose dendrites serve the rods may derive from many of these receptors and is thus more likely to pick up slight changes in intensity of the light striking some part of the retina. Peripheral to the fovea, however, as has been shown by Polyak, some of the bipolar cells synapse with both rods and cones so that the duplex nature of these systems is not absolute.

a. Student should understand the difference in the innervation of the two types of end-organs. (*Rods do not have individual innervation; cones in the centre of the fovea do.*)

b. Student should understand the implication of the differences in innervation for acuity.

c. Student should be able to explain why the duplex nature of the system is not absolute.